Picture credit: Shutterstock.com

INDICATORS OF BEHAVIOUR

Sam Curry *and* **Anthony Freed** *offer up an orthogonal approach to early attack detection*

Despite all of the tremendous progress the security industry has made over the last 30-plus years, actually solving one critical problem seems to remain as elusive as ever: how do we detect and stop advanced attacks before they escalate to full-scale security events?

To answer that question, we will need to do some things that are, by human nature, difficult; namely, we will have to try something different. Why something different? Because we keep trying to leverage Indicators of Compromise, derived from known attacks in one environment, to proactively detect new attacks in other environments.

This article outlines some of the reasons for the development of an extensible language to both detect and describe the most subtle chains of malicious activity that will result in better detections earlier in the attack kill chain, a concept that will hereon be referred to as *Indicators of Behaviour (IoBs)*.

It's important to hire the right people with the right skills and to move resources to functions that are more important because they work

THE DIMINISHED VALUE OF INDICATORS OF COMPROMISE

The fact is that Indicators of Compromise (IoCs) are constantly changing and typically unique to a specific target, so leveraging them for proactive defence is a challenge. Even the assumption that IoCs are somehow uniformly applicable in every instance for a given attack campaign in the same environment is demonstrably false. And to further complicate the issue, attackers even change techniques and tools within the same kill chain from one device to the next.

Furthermore, the attackers are quite aware that we in the security community readily share this intelligence, so when analysts search for IoCs in public repositories like VirusTotal, the attackers can see the queries and we're effectively telegraphing the techniques and tools that have been picked up on and those that are still effective. This intelligence feedback is gold for the attackers, and they are consuming it and adjusting their TTPs accordingly. Checking IoCs may actually be more of a liability against sophisticated attacks and operations in this case, and what's sophisticated today is usually commoditised tomorrow.

And of course, we still haven't found a solution to the supply chain issues that made for big headlines in late 2020 and early 2021. This creates situations where we are trusting a piece of software in an environment and assuming all the security was done right, yet bad things can still happen. Especially when those tools can be used against us, as we saw with SolarWinds software and Microsoft Exchange software around SUNBURST and HAFNIUM, respectively.

RELIABLE DETECTION

So, how can we detect more reliably and earlier in the kill chain? We need to solve the issue of noise-to-signal ratio first. Most of security for the last two decades has been based on the idea that if we just collect enough security telemetry in one place, we can just sift through it and find evidence of an attack.

This notion is based on something called Locard's Principle from the 19th century. Locard was a French criminologist – if you've ever seen a CSI show, you'll know what this is. It's the idea that whenever a criminal interacts with a crime scene, there's an exchange of evidence between them and that crime scene. And if we can just freeze the scene and find that interaction evidence, we can reconstruct what happened during the crime. And while that may be true, the principle has yet to be proven to be applicable for cybersecurity forensic examinations.

So, here's our challenge: put simply, we need to catch the attackers closer to real-time and earlier in the attack sequence – and for the most advanced attacks, we are not going to be able to do so effectively based on IoCs from already realised attacks. Worse yet, there may be canaries in the coal mine for attackers to gain an edge in offence, and therefore IoCs really shouldn't be used if possible. That means we need to focus on collecting and enriching the right telemetry, and assuring a low noise-to-signal ratio. And we have to be able to stop supply chain attacks – in particular, we have to be able to say that even when trusted software goes wrong, we can still catch it early.

We need the ability to detect advanced techniques used for initial ingress, to establish persistence, to elevate privileges, to compromise user identities and to quietly move through a network long before the actual payload is delivered. We need something to find all of this and operationalise it in defence of our networks, and it needs to be something we can share with other security practitioners in the way we share IoCs.

Thus, we are proposing the development of an extensible language that effectively detects and describes the most subtle chains of malicious activity derived from enriched telemetry from across all network assets, intelligence we will refer to as Indicators of Behaviour.

WE NEED TO CATCH THE ATTACKERS CLOSER TO REAL-TIME AND EARLIER IN THE ATTACK SEQUENCE

DEFINING INDICATORS OF BEHAVIOUR

Unlike retroactive IoCs, Indicators of Behaviour is a proactive approach to leveraging real-time telemetry, and it is our intention that a science and new technologies can be built around them that will enable us to have a more future-proof approach to detecting novel and emerging threats - or at the very least a system for finding attackers that has more longevity than previous techniques have been able to deliver.

David Bianco, many years ago, came up with something called the Pyramid of Pain, which is a hierarchy of telemetry. At the base, we have things that are easy to get and that we've used for years like hashes. And then going up the pyramid, we have things like IP addresses and domain names. Now, these are getting progressively more difficult to ascertain as we go up the pyramid, so by the time we reach the top part of the pyramid we're getting to the stuff that's really hard to ascertain, but which also has the most longevity and value.

Getting to the very top is hard – that would be identification of previously unknown or unrelated techniques, tactics and procedures (TTPs) – all the activities and actions telemetry for the attackers on the network. Being able to instrument those TTPs by correlating and contextualising them to make them actionable is at the heart of being able to leverage IoBs, which are essentially specific to the telemetry for TTPs in use.

In the world of SIEM (Security Information and Event Management), we aggregate and record all sorts of security telemetry, but if no standing policy is actually triggered, then all that telemetry that we could leverage to find similarly novel attacks in the future is for the most part completely useless.

The ideal state would be to shrink the budget spend at the bottom of the pyramid where it has become bloated yet ineffectual, then invert that pyramid entirely. We are not advocating we simply

▶

throw this stuff out – this intelligence is still quite useful, and we should and will continue to use it all. But that doesn't mean that it should be at the centre of our detection strategy. Ultimately this will free up resources for what we really need to level-up our security programmes to match the threat from the adversary, to make sure we can hire the right people with the right skills and to move resources to functions that are more important because they actually work.

> ATTACKERS OFTEN CHANGE TECHNIQUES WITHIN THE SAME KILL CHAIN FROM ONE DEVICE TO THE NEXT

So, what are we really looking for? Let's take a queue from the MITRE ATT&CK. What is so valuable about the MITRE ATT&CK framework is it provides a taxonomy for us to describe the art of what the bad guys do. What do the attackers do when they move through an environment? The columns display the tactics throughout the stages of the attack, and the individual boxes are the techniques that they use to accomplish each of them, and what we're trying to do is to find these and the links among them on our network.

These are the trajectories, pathways and sequences that will stand out from the background noise if we had a way to effectively uncover them by leveraging correlations across the most meaningful telemetry to detect a particular moment in the kill chain. This isn't about finding one of these tactics, it's instead an orthogonal view of an attack where we're detecting the pathways through these techniques.

Every box in the MITRE ATT&CK represents a chance to find attack activity based on subtle behaviours, and then to build from the chains of behaviours that have come before it. Detections based on chains of behaviour – even those that are normally benign that we'd expect to see on a network – become suspect because they are either rare or present a distinct advantage to an attacker, allowing defenders to respond in real-time to stop the activity. Our goal, again, is to find things more reliably and sooner, pushing detections further to the left on the attack timeline.

SETTING A STANDARD

It's time for a future-proof standard to define and operationalise Indicators of Behaviour so we can reliably and repeatably reveal the earliest stages of an attack – to more rapidly get to the actual DNA of the actions and the activities of the attackers. And it's all independent of which security tools happen to be in place; they can provide the colour, they can provide the context, but the tools don't provide a language that actually describes these chains of behaviour and let's us respond to them faster.

We need a common, extensible format for IoBs that can keep us all on the same page yet is capable of scaling as our capabilities and those of our adversaries continue to evolve. This does not mean simple anomaly detection, or shifting our focus to some sort of behavioural anomaly detection, or launching some sort of UEBA redux.

Instead, what we're talking about is instrumenting and collecting behaviour at scale – both good and bad – and putting it into data structures that can enable queries while allowing additional context from diverse telemetry sources. The concept of IoBs in practice will filter the noise out, but keep the things that matter, and produce a standard for this future-proof telemetry that will collectively benefit everyone going forward.

To that end, we are working with other security professionals through OASIS to take the next steps in developing the foundation for a common IoBs Standard, and we welcome all to participate (join the mailing list at oca-iob-wg+subscribe@lists.oasis-open-projects.org). It's time to start thinking about how we leverage Indicators of Behaviour in addition to the more common Indicators of Compromise as the primary, proactive approach to detecting attack activity and actions as early as possible ●

Sam Curry is Chief Security Officer at Cybereason and **Anthony Freed** is Senior Director Corporate Communications.

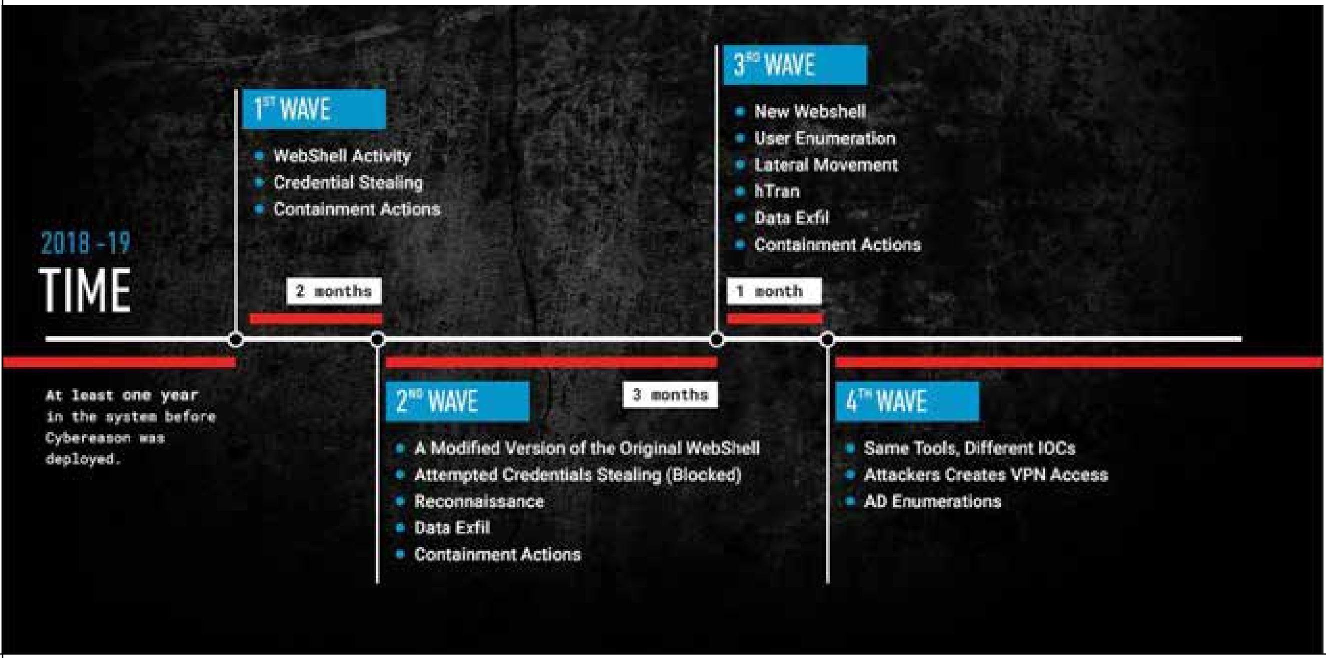

2018 -19

TIME

At least one year in the system before Cybereason was deployed.

1ST WAVE
- WebShell Activity
- Credential Stealing
- Containment Actions

2 months

2ND WAVE
- A Modified Version of the Original WebShell
- Attempted Credentials Stealing (Blocked)
- Reconnaissance
- Data Exfil
- Containment Actions

3 months

3RD WAVE
- New Webshell
- User Enumeration
- Lateral Movement
- hTran
- Data Exfil
- Containment Actions

1 month

4TH WAVE
- Same Tools, Different IOCs
- Attackers Creates VPN Access
- AD Enumerations

Low and slow: the mark of a persistent attacker

Picture credit: Cybereason