CYBER STASH

Anatomy Of A Cyber Attack

Understanding The Role of Defensive Technologies and Forensic Depth Analysis in Breach Detection and Prevention

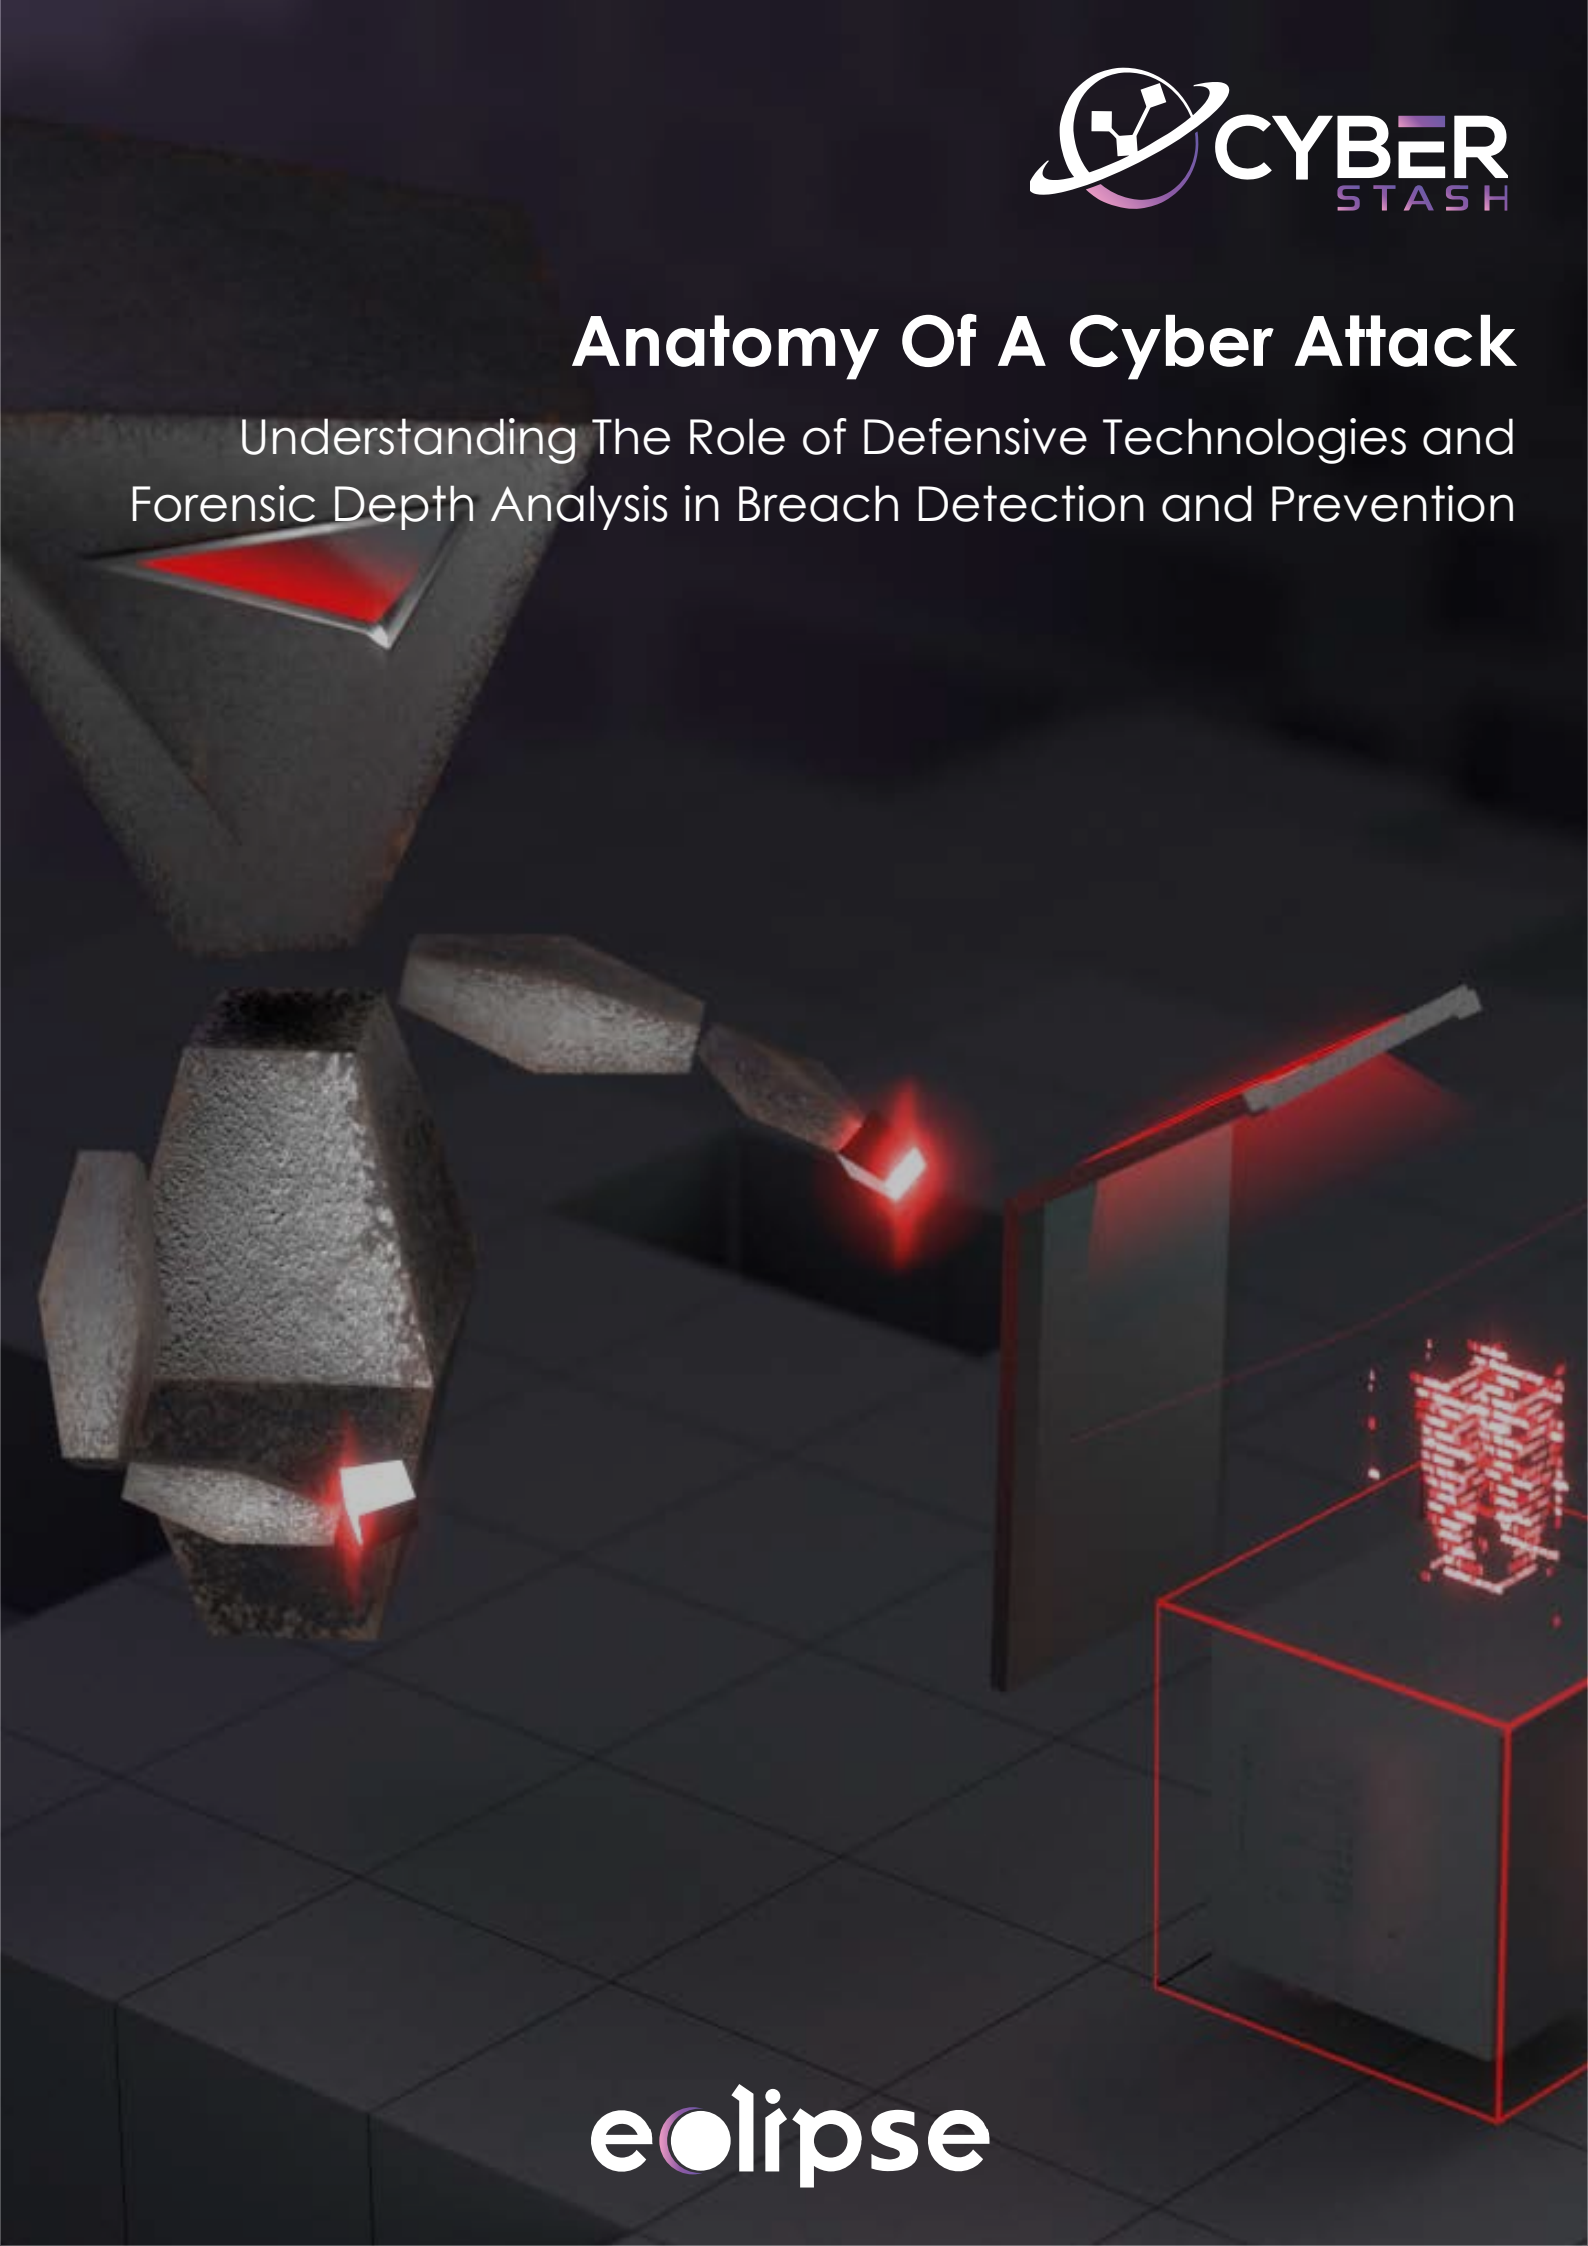

eclipse

Table Of Contents

Executive Summary
The Setting
Act 1 – Attacker Entry
Endpoint Detection and Response
Act 2 – The Attacker Is In. But How Would You Know?
Forensic Depth Analysis
Act 3 – We Found Our Attacker. Now What?
Digital Forensics Incident Response
The End. What Did You Learn?
FDA is the Hero
Parting Thoughts

Executive Summary

Threat Hunting is the search for unknown compromises and threats that have already bypassed prevention-oriented security controls.

More than just hype, threat hunting is a legitimate and necessary tactic for modern cybersecurity practitioners. The top efficiency benefits from a threat hunting platform are:

- ***Improving the detection of advanced threats***
- ***Creating an independent method for discovering threats***
- ***Discovering threats that could not be discovered otherwise, and reducing investigation time.***

The benefits of threat hunting impact the bottom line because how quickly an organisation contains a data breach has a direct effect on the financial impact. Case in point, the impact of a data breach is reduced by 96% for organisations that were able to contain the breach in less than 1 day.

Looking to capitalise on the benefits, the security market has suddenly become crowded with solutions that all claim to offer threat hunting capabilities: EDR, DFIR, Behaviour Analysis and FDA.

Understanding the differences between threat hunting tools and the role each plays in breach detection and prevention

Threat hunting with FDA or Forensic Depth Analysis offers a unique approach that is complementary to other threat hunting approaches. It is not a replacement for alternative approaches like Endpoint Detection and Response (EDR) or Digital Forensics and Incident Response (DFIR). That said, FDA provides the most conclusive post compromise detection ability, is the easiest to use, and is by far the most cost-effective approach on the market.

The purpose of this paper is to explain FDA in more detail, such that hunt practitioners, security budget decision makers, and risk management leaders can understand why deep memory state analysis provides so much promise in the fight to stop adversaries from reaching their ultimate theft or damage objectives. It also introduces CyberStash Compromise Assessment Service, a threat hunting solution that offers post breach detection using Forensic Depth Analysis (FDA) to discover hidden threats and compromises within a network.

The Setting

You've seen this play out in a Hollywood movie more than once. The one where there is a jewel heist from within a supposedly well-guarded building. Except today's thefts have moved online where the stakes and payloads are higher; company data, credit card numbers, personal data, corporate IP, and even access to critical infrastructure.

The Scene

The antagonist. There is a hacker. He wants in. But he knows there are lines of defenses. So he must understand them, and figure out how to evade them, at least until after the "jewel" is in hand and he is well out the door.

The protagonists. The security folks, IT team, and managed services providers whose job it is to prevent the hacker's entry, find him if he trips a wire along his journey, or investigate how he did what he did, if he somehow escapes with the jewel – so we can maybe we can stop him in his tracks, but at least learn from our "defense in depth" shortfalls.

Well, there is (arguably) a rough analogy here – but one that is good enough for our purposes. And, that analogy is to break down the anatomy of the attack into movie acts – so we can see where different "hunt tools" can be brought to bear.

Act 1
Attacker Entry

The first step is to get into the building. But there are system alarms, security guards mat the front desk, security cameras, maybe motion sensors in a few hallways, elevator movement monitors, etc., Action! Let's catch the attacker in the initial act – an actual breach in process!

Enter EDR.

Endpoint Detection And Response

EDR products typically rely upon behavior monitoring and analysis technology for their purported "hunt" capability. Their approach is to record changes to a system (or network) as events (new process spawn, registry key change, or user privilege escalation) occur. Examples of recorded data include:

- Process execution events (occasionally with command line used)
- Process changes (elevation of privileges, process crashes, etc.)
- Select registry changes / writes
- Select disk writes, i.e. download/user folders, windows folder, etc.
- File creation events
- Monitoring select API calls (monitoring all would be impossible)
- Sampling of network connection events

These are valid things to monitor – provided the goal is to catch an attack in progress.

BUT what if our hacker is especially crafty? What if he is equipped with evasion techniques that render your behavior monitoring defense-in-depth systems inert?

You've seen this:

- Gas that puts the guards to sleep.
- Jumper wires that prevent alarms from activating.
- Still photos in front of the monitoring cameras.
- Mirrors that trick laser trip wires.
- An un-monitored entry point. One that you didn't know could be breached.

Oops…sounds like a zero-day!

Now we have a problem. The hacker is in the network. What about all those front-line defenses set up to detect and alert on entry? Well, our guy is behind them now so they're kind of, to be blunt, useless.

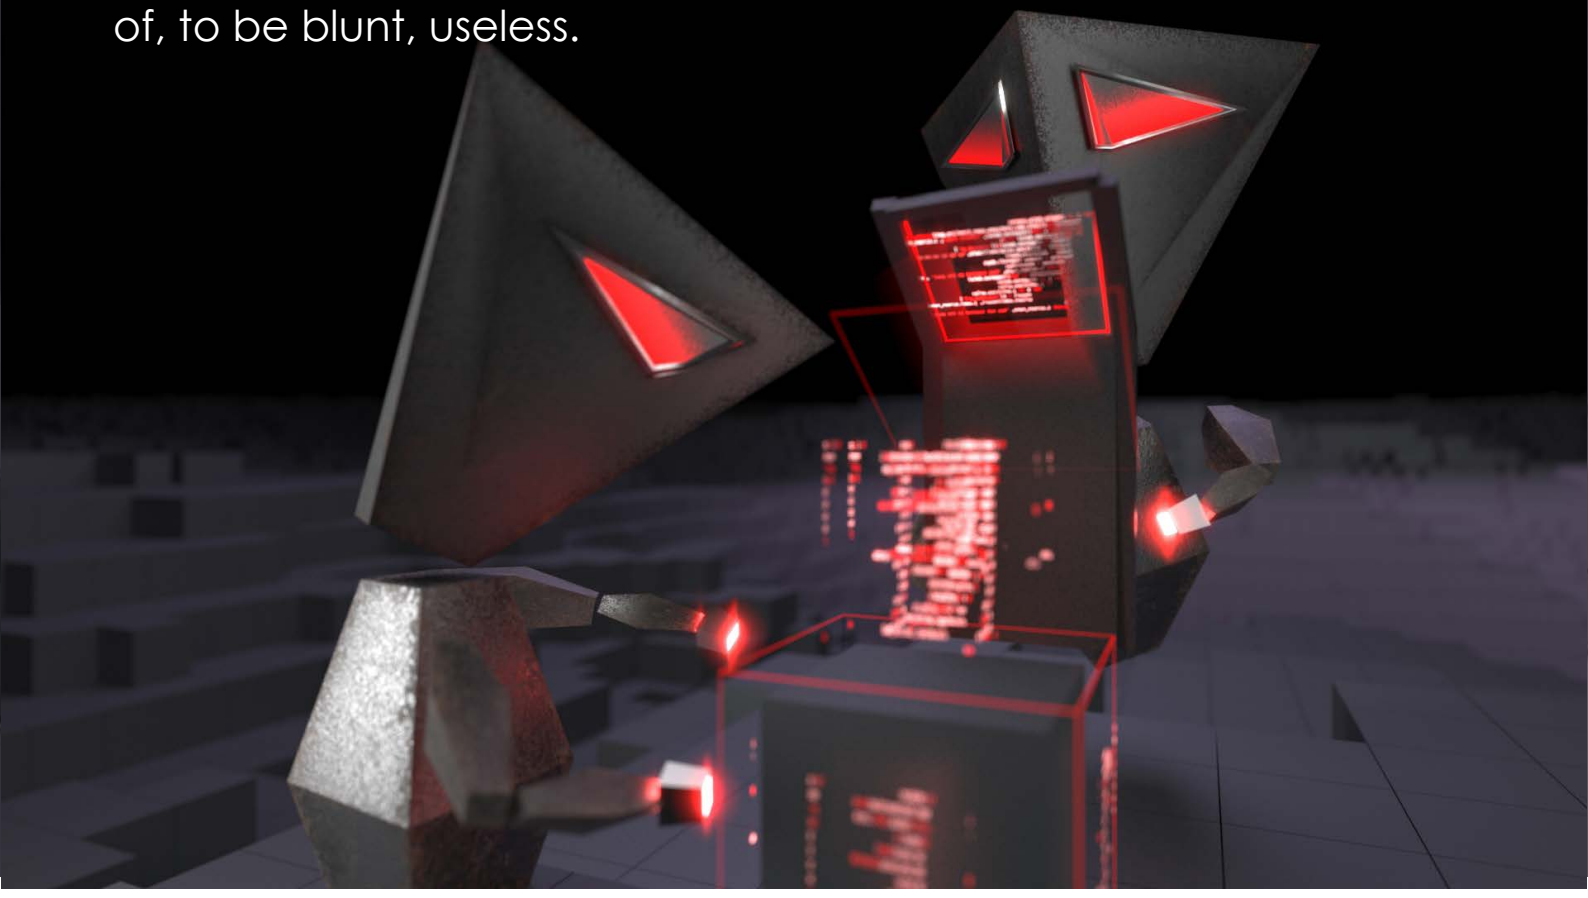

Act 2
The Attacker Is In.
But How Would You Know?

Our guy is in. Now his goal is to sneak around until he finds the exact right room; server, network, PC. The one with the jewel in the glass case, the payload. But this will take time. And so he must move intelligently, staying under cover and never dedicated. And as he moves, he'll be searching to find keys – user passwords, admin credentials, etc. – in one 'room' that will help him enter the next (without being noticed), if he can find the right disguise.

Hold on. What if we had systematic monitoring of every room – or endpoints in this case? And, what if that monitoring was super sophisticated, able to see individual footprints; suspicious code, memory injected modules, process manipulations, etc. – all on a per endpoint basis? Sounds great, but that's a lot of data to collect and analyze. But if we could do it, we'd have a great way to detect the whereabouts of the attacker who has compromised our frontline defenses.

This is where Forensic Depth Analysis (FDA) enters the picture.

Forensic Depth Analysis

At the highest level, FDA assesses three things in detail:

1. What is actively running on an endpoint
2. What is triggered to run – through a persistence mechanism – on an endpoint
3. The identification of any operating system (OS) manipulation, or active process, e.g., what a rootkit does to hide its presence, or what an insider threat might do to disable the system's security controls

Examples of findings include things like unusual OS configuration settings, or API calls being hooked by a rogue/hidden process within volatile memory, i.e., a rootkit.

FDA does not rely on logs or monitoring events/changes to a system. FDA assumes the device is already compromised and seeks to validate every aspect of the system as deeply as possible. CyberStash Compromise Assessment Service uses FDA to discover hidden threats and compromises. It sweeps thousands of endpoints, spending a couple minutes on each host, and conclusively validates their state: "Compromised" or "Not Compromised". To accomplish that, Forensic Depth Analysis takes 13 Steps to definitively establish trust in an endpoint.

13 Steps For Conclusive Validation

1. Evaluation of all active processes.
2. Evaluation of all loaded modules and drivers.
3. Identification and evaluation of all memory injected modules.
4. Conduct memory un-mapping techniques – which are used to export memory objects for offline retention and analysis.
5. Identification and evaluation of process manipulations, e.g., function hooks and in-line modifications / patches.
6. Identification and evaluation of operating system manipulation including list modifications, hidden parocesses, and direct kernel object manipulations.
7. Identification of disabled security controls, e.g., disabled anti-virus, reduced authentication requirement configurations, GPO blocking.
8. Enumeration and evaluation of persistence including cron-jobs, registry auto-starts / triggers, DLL hijacking, WMI Events, boot process redirection and watchdog processes.
9. Evaluation of application execution artifacts, e.g., Prefetch, Shimcache, and SuperFetch.
10. Identification and evaluation of web shells – Linux or IIS web servers.
11. Auditing of legitimate remote admin services like cmd, Powershell, NetSH, SSH, VNC, PSExec, RDP, Tunnels and WMI.
12. Evaluation of all active host connections, including inter-process and redirects.
13. Auditing of all privileged user accounts, e.g., ID rogue local admin accounts.

Bypassing Anti-Forensic Techniques

To establish trust in endpoints, successful state analysis also requires the ability to bypass anti-forensics techniques. CyberStash Compromise Assessment Service accomplishes this by:

- Going underneath higher-level operating system APIs
- Working directly with volatile memory structures using patented memory analysis techniques.

So you see, post compromise detection is different from finding an attack in progress as it is completely independent from the Operating System or the method that was able to compromise the system in the first place.

Our prospects often ask,
"How does the CyberStash Compromise Assessment Service perform behavior analysis if it is agentless?"

Well, it doesn't!

With exception of sandboxing during binary analysis phases, CyberStash does not use behaviour detection techniques at all.

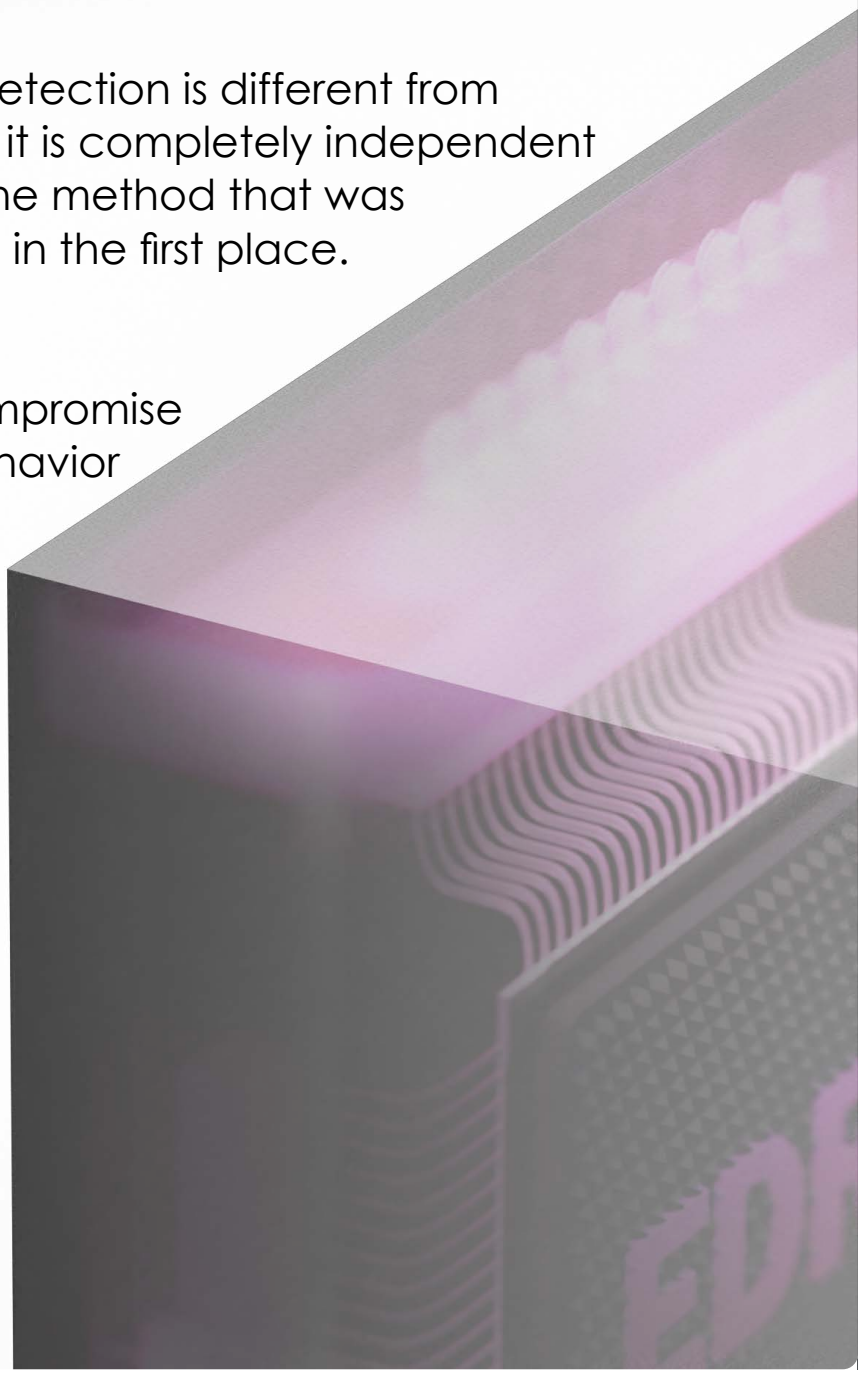

CyberStash Forensic-Depth Analysis

Given the extensive endpoint state collection and analysis performed by CyberStash, it comes as close as possible to being able to assert 'this endpoint is clean'. EDR will never be able to make this claim. It is simply not designed to do so. Again, EDR tools monitor endpoints for behaviour indicating an attack is underway. They do not perform forensic validation of endpoint cleanliness.

To further drive the point, let's imagine you were asked to ensure a house was clean from intruders.

EDR, and its behaviour monitoring feature set, is centred on the premise that if you monitor all doors and windows, no unwanted person could possibly be inside the house. Breach after breach has clearly proven this to be patently false. Now, no one is suggesting that FDA is somehow perfect. However, it should be clear that analysing the state-of-affairs of a room past the point of entry delivers unique and compelling value for post-breach detection.

Here's an example that fits our story.

The security personnel at the hotel can be the most sophisticated in the industry, have all the latest locks on the doors, security cameras, and exterior defenses a hotel could possibly use (analogous to EDR). And yet, the secret service is still going to sweep the room for bugs and verify no intruders are currently in the room (FDA). A room swept for bugs, using a reasonably comprehensive process, is far safer than an un-swept room. The FDA sweep is completed independent of whether or not any of the existing in-line controls detected an intruder.

It provides assurance.

Act 3
We Found Our Attacker. Now What?

There are only three ways this can end:

1 – Our hacker gets in and gets out with the jewels, and, we never saw him.

In fact, we are so weak at security, we don't even recognise the jewel is gone until six months after the fact.

Sadly, many corporate cyber security stories play out this way.

Why?

Well first, they trusted defense-in-depth. Second, in a corporation, there isn't just one jewel. There are hundreds to thousands of them. Intellectual property repositories. Credentials left and right. Personal privacy records by the truckload. Payment card data. In fact, there are so many, that we don't even know where they all are at a given point in time.

2 – We catch our hacker as he enters, i.e., while his entry attack is underway. Great!

He got past our prevention investment which is where we spend most of our budget, wrongly, given modern burglar sophistication. But thankfully, he tripped a behavioral wire. We detected this "indicator of compromise". And then, we had our really bright security analyst (who is rare, busy, and very expensive) take a look at this indicator, match it up to a lot of other big data, and conclude – correctly (this time) that it is in fact a real burglar, and not the janitor

3 – We catch our attacker – who cleverly got past all our prevention and entry detection mechanisms – but was not able to evade our FDA hunting methods.

In any of these outcomes, the third act is all about detaining him if we can, and then once the dust settles, figuring out how he got that far – so we have a clue about how to prevent that movie from going to a sequel.

This is where Digital Forensics Incident Response (DFIR) comes in.

Digital Forensic Incident Response

Digital Forensics and Incident Response (DFIR) is the procedure of investigating security alerts or suspicions of malicious activity in a computer network. Enter EDR.

By examining a breach or an attacker's infiltration in detail, a skilled forensic analyst can come to understand misconfigurations, lack of security measures that might have allowed the attack to take place, and attack details that can assist remediation.

DFIR solutions are fabulous. It's the blue light in the forensic specialist's hand that finds all kinds of nasty evidence at the scene of the crime. It goes down to the granular level of DNA analysis of a hair.

There is just one problem. DFIR is designed to focus on the deep data collection and analysis of a single endpoint at a time – and by yet another highly skilled, highly specialized, and, therefore, very expensive analyst. Let's face it, you just cannot afford to perform DFIR on every endpoint at a high enough frequency to detect an attack in progress, let alone the more difficult work of sweeping an entire building for post compromise detection.

DFIR, while extremely effective in providing conclusive evidence of compromise, is simply not a technique that we can adapt for advanced threat detection at mass scale.

The End.
What Did You Learn?

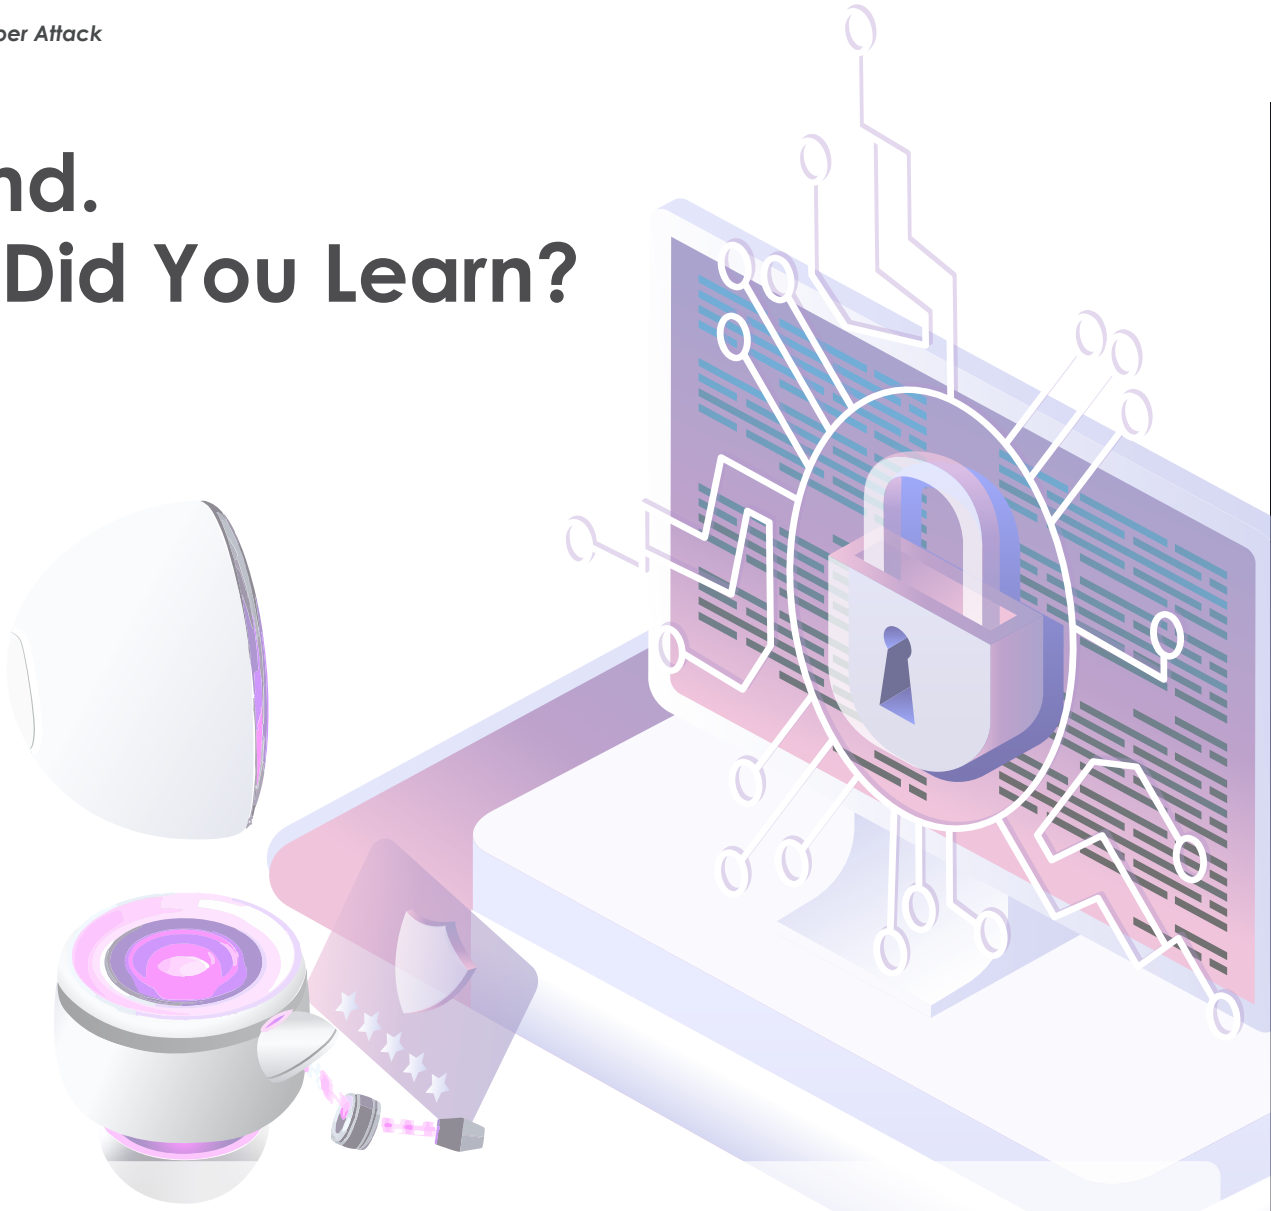

In summary

Act 1. is straightforward. Stop him from entering via prevention. Or detect his entry via indicator of compromise and/or behavioral analysis. That can work. But, as we know, often it does not. So this act is not that "intriguing".

Act 2. He got in. He is smart. He is crafty. He knows when to move and when to lay low. Ah, now we have an interesting story. And it's game on but only if you are using the CyberStash Compromise Assessment Service with FDA.

Act 3. Frankly, boring. Yeah, it's cool to know how the guy did what he did, but the damage is done if he succeeded.

FDA Is The Hero

The hero of our story is FDA.

FDA is not a replacement for centralised logging or real-time behavior monitoring because these less than effective discovery approaches have their place in helping organisations to comply with regulatory requirements.

For mature enterprise SOCs and organisations who take post-breach detection seriously:

FDA enables the elimination of custom scripts and/or single-host-at-a-time DFIR processes used to validate suspicious behaviors detected by your team

FDA enables hunt teams to iteratively and effectively sweep all endpoints to find entrenched threats and beachheads hiding on any of your endpoints

Now, many advanced SOCs are likely already doing a lighter version of the above, but with a custom tool set or scripting out an endpoint-querying tool. Not only are these approaches difficult to scale and maintain, they will have limited effectiveness – as they are unable to bypass anti-forensics safeguards.

FDA Is The Hero

For newer moviegoers (i.e. newer or smaller Hunt teams), CyberStash provides, by far, the biggest bang for the buck. CyberStash incorporates the FDA methodology to automate most of the workflow and analysis for you. We find relatively conclusive results. So you do not require a department full of hard to hire, hard to retain, expensive security specialists.

Finally, whether you have a sophisticated SOC, or are just at the early stages of learning to hunt, CyberStash not only supercharges your monitoring and threat hunting processes, it enables entirely new use cases:

- Laptops, mobile devices, and other transient systems not previously under management can now be validated as they join the network.
- Systems without endpoint monitoring (due to policy, mismanagement, or tampering) can be identified and periodically assessed.
- Organisations that don't have sufficient historical log data, or the ability to convert big data into definitive action, realise huge value from FDA
- Consultants and IR professionals now have access to the fastest and easiest way to perform a compromise assessment or threat hunting engagement service using CyberStash.

Parting Thoughts

Movies are fun. We all love them. Cyber attacks are not fun. They will cost you money, your company's reputation – even your job. You don't have time to waste on rabbit hole analyses. You don't have enough staff. You'll never have enough budget.

CyberStash Compromise Assessment Service is the industry's leading Forensic Depth Analysis platform. No enterprise should be without it. Security is a tough game. And certainly, other tools have their place. But CyberStash with its ability to perform post breach detection service is your best possible threat hunting investment.

info@cyberstash.com
1300 893 802
cyberstash.com
Sydney, Australia

Want to run a trial?

Reach Out To Cyberstash
For More information.

Explore
eclipse.xdr

CYBER STASH