CYBER STASH

Post-Breach Forensic-Depth Compromise Assesments

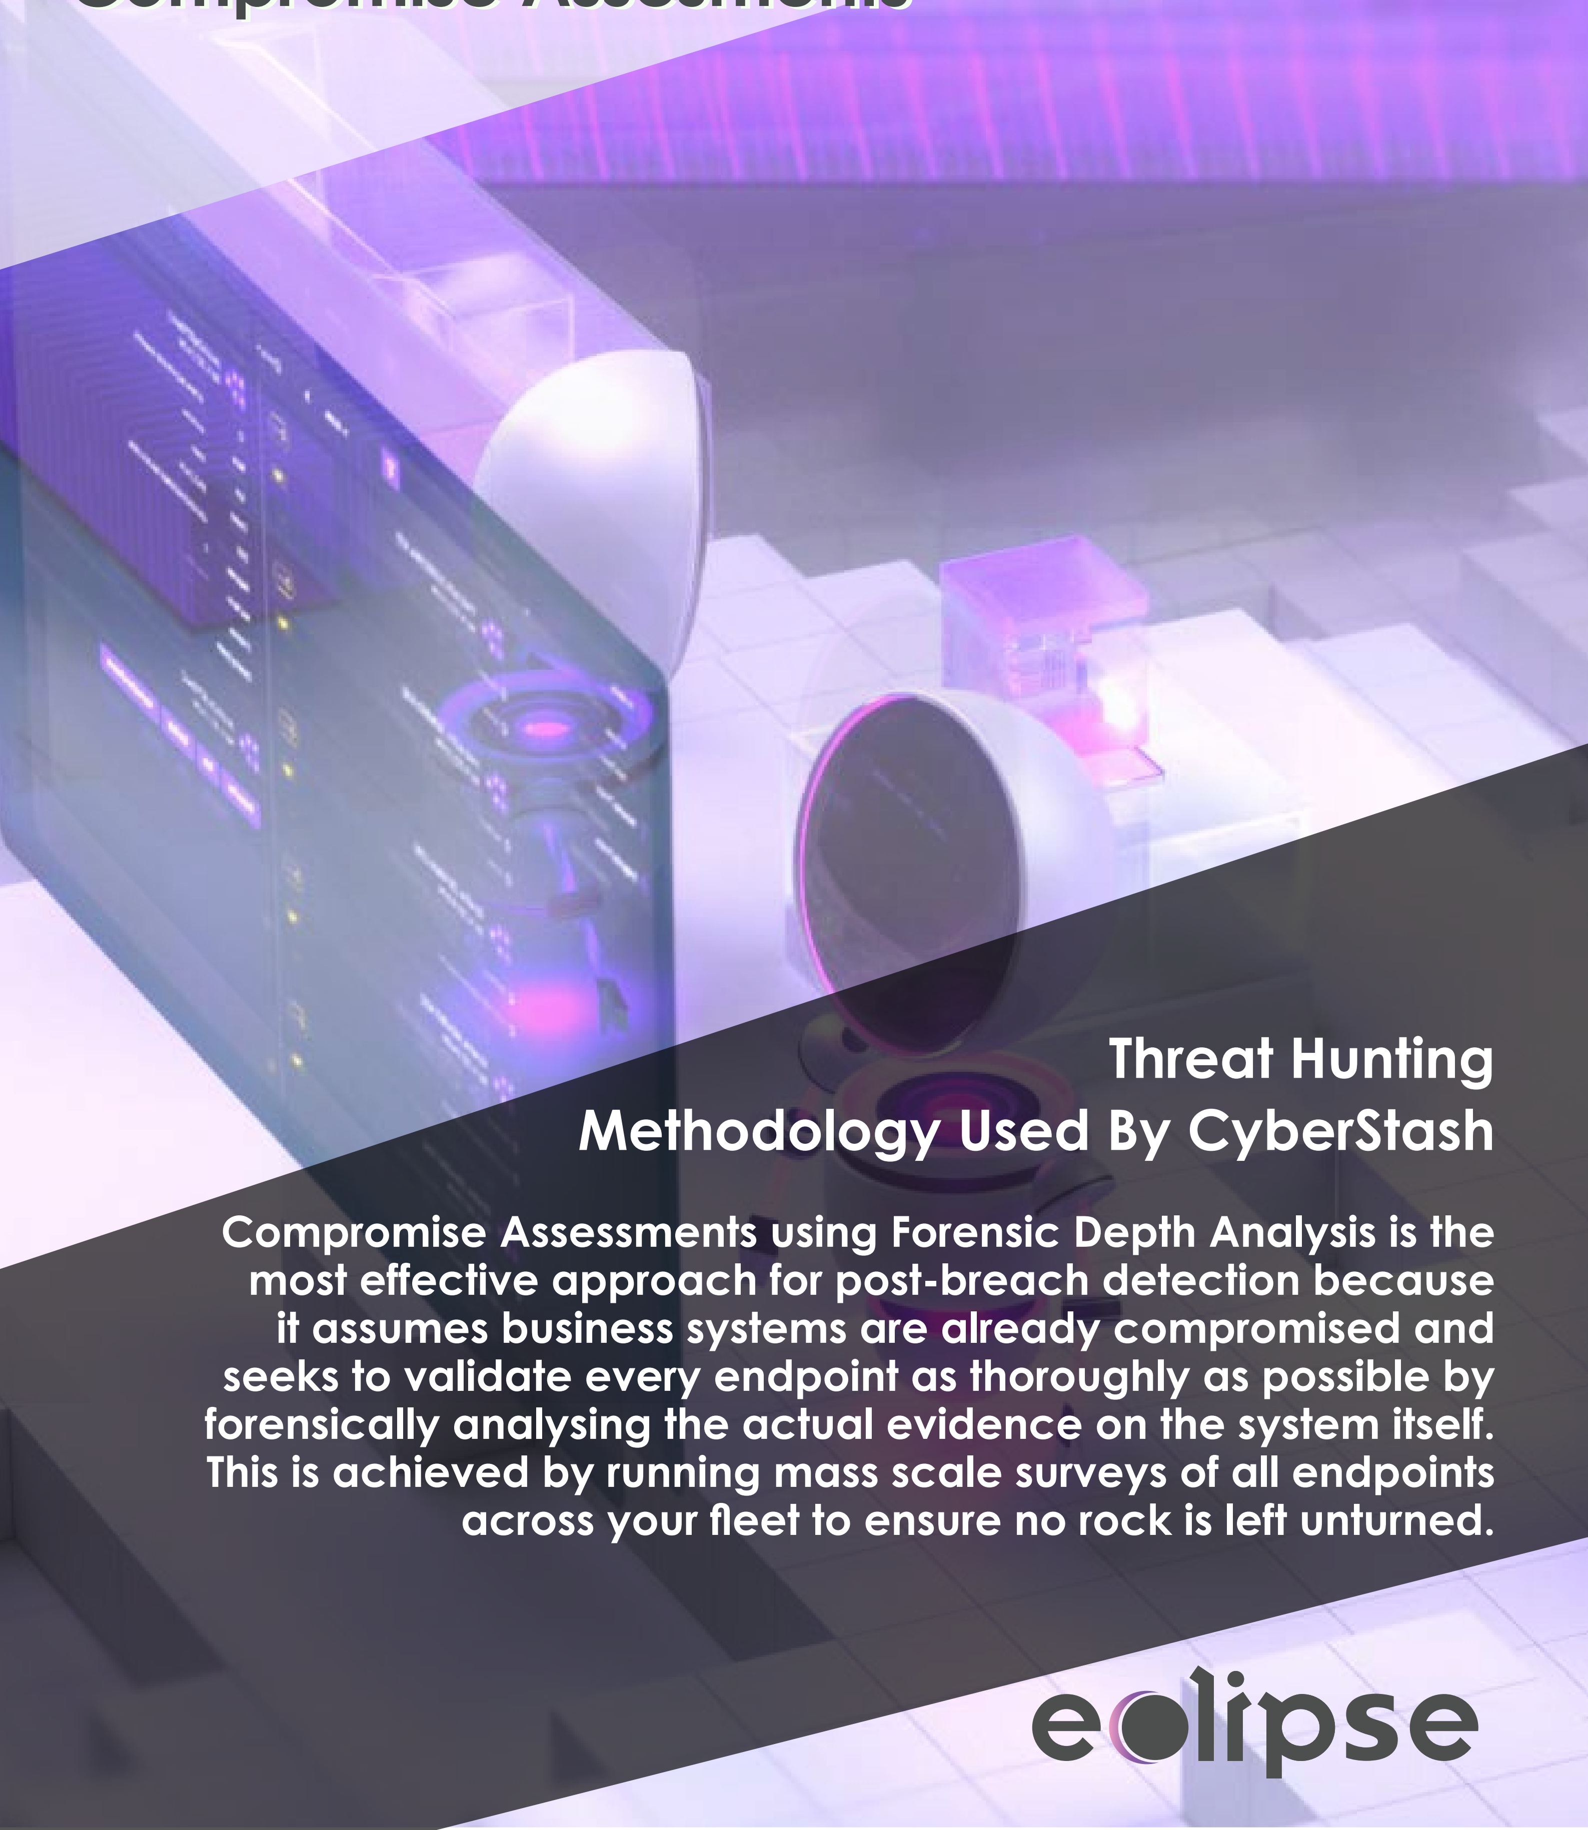

Threat Hunting Methodology Used By CyberStash

Compromise Assessments using Forensic Depth Analysis is the most effective approach for post-breach detection because it assumes business systems are already compromised and seeks to validate every endpoint as thoroughly as possible by forensically analysing the actual evidence on the system itself. This is achieved by running mass scale surveys of all endpoints across your fleet to ensure no rock is left unturned.

eclipse

Context

A World Without Certainty

Even organisations with the most effective security practices fall victim to advanced and persistent threats. With sophistication and elegance, hackers compromise systems and information while remaining undetected.

The advancing nature of threats, means our current reality is one where risk, compliance and security practitioners are unable to provide assurance to business stakeholders.

While the practice of penetration testing, vulnerability management, and red/blue teaming provide us a hackers view into our networks and tell us what our weakest links are, they fail to confirm whether any of our business systems are already compromised.

The question to then ask is this:

"If the results from a recent penetration test demonstrated that it's in fact possible to breach your organisation, what makes you think that the breach has not already occurred?"

To change the economics of cyber defence so that the defenders have the upper hand, we must begin to think like the adversary, mirror their every move and hunt them down where they live, which is on our endpoints, and, we must achieve this while optimising risk and resources.

Detection Strategy

CyberStash's Compromise Assessment Service utilises forensic depth analysis (FDA) to perform deep host inspections of devices. Deployment models can be either agentless based, agent based, or both. Unlike other analytics solutions that focus on behaviour (e.g., UEBA), CyberStash collects its own primary forensic data rather than relying on existing security logs from sensors (IDS, AV, etc.) that might have failed to alert on the attack in the first place.

The key premise behind CyberStash's approach is that a performing log analysis — the key method used by most organisations — is generally expensive, difficult to manage, and error prone. Log analysis approaches require in-depth knowledge of adversary tactics and how those tactics present themselves in the logs of security solutions. Log analysis typically requires that a product (and security analysts) performs a great deal of tuning around the exact devices and information that the solution collects, as well as then determining if devices are reporting the correct information to begin with.

CyberStash's solution and service was designed with many principles in mind, namely independence, minimal invasiveness, and simplicity. CyberStash begins by assuming that endpoints are already compromised and seeks to validate that assumption using a variety of forensic and threat hunting techniques. Automated forensic collection, volatile memory inspection, threat intel enrichment, and deep analysis workflows to dig into anomalies and outliers help hunters find what purely automated detection misses.

Using FDA To Hunt For Persistent Compromises

CyberStash uses FDA to discover hidden threats and compromises on endpoints. It continuously inspects endpoints at scale, spending a couple minutes on each host, and seeks to validate their depth as "compromised" or "not compromised.

At the highest level, CyberStash digs deep into endpoints to validate:

What applications and processes are running (in memory)

What is triggered to run (through a persistence mechanism) or has run previously (via forensic execution artefacts like Microsoft Windows Shimcache)

It identifies any manipulation of the operating system (OS) or active processes. This will reveal things like an OS configuration setting, or an API call being hooked by a rogue/hidden process within volatile memory (e.g., rootkit).

This process of hunting with FDA is performed in five steps:

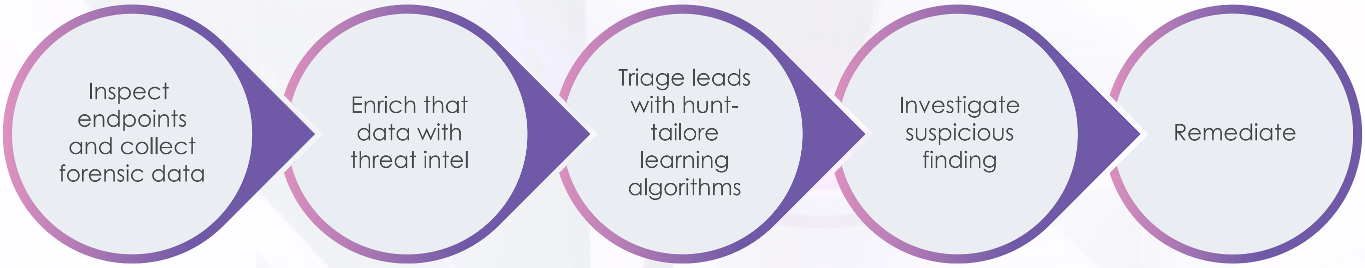

Inspect endpoints and collect forensic data

Enrich that data with threat intel

Triage leads with hunt-tailore learning algorithms

Investigate suspicious finding

Remediate

This is a highly differentiated approach from the behaviour analysis techniques used by endpoint detection and response (EDR) or user behaviour analytics (UBA) products, which only record the changes to a system or network as events (e.g., a new process spawning, a registry key change, or a user elevating privileges). FDA digs much deeper into each host. Perhaps the most important aspect of ensuring that the depth analysis of a compromised machine is successful is being able to bypass anti-forensics techniques. This is accomplished by going underneath higher-level operating system APIs and working directly with volatile memory structures — both of which CyberStash does.

Endpoint Focused Threat Hunting

Threat Hunting

Threat hunting focuses on proactively and systematically searching through data points to detect and isolate threats that have evaded existing security controls, bypassing defences and are residing undetected. Today's hunting solutions generally focus on working with data points available either from networks or endpoints.

Endpoint Focused

Focusing attention on endpoints is a logical and more effective approach. When an adversary is embedded in a network, they will most likely conceal themselves on an endpoint to utilise it as a launchpad for attacks.

Endpoint focused threat detection encompasses several types of solutions, namely: Endpoint Detection and Response (EDR), Security Intelligence (SI), and Forensic Depth Analysis (FDA).

Both EDR and SI fall under one umbrella, they rely on log collection, pattern-matching and behaviour analysis. Forensic Depth Analysis, by contrast, is quite different. Let's first

Problems With Log Based Threat Hunting Approaches

Collection

The National Institute of Standards and Technology (NIST) Cybersecurity Framework and similar frameworks, depth that apart from protection, organisations should equally focus on their detection and response capability. Unfortunately, this has been translated by many to mean the collection and analysis of second-hand information collected from network, application and system events. While using a Security Information and Event Management (SIEM) system may have been our initial answer for building detection and response capabilities, the efficiency and effectiveness of these solutions no longer justify their total cost of ownership. Collecting network, application and system events, demands a large footprint, both from system resources and human resources. This is a costly exercise and one that returns minimum results by way of leading to the detection of advanced threats.

Intelligence Gaps

There's over 950 million threat indicators freely available from the open-source community alone. Harvesting this intelligence and operationalising it to make it actionable is a massive security challenge and one that requires dedicated resources. Not having timely access to the right intelligence equates to having security gaps.

Correlation

Even when you do have timely access to intelligence and you have collected all the network, system, and application logs, organisations must then correlate these two datasets in close to real-time as possible and attempt to detect threats that have circumvented existing controls. Processing power is then required to categorise and correlate relevant indicators with the relevant pieces of logs, making the log-based approach to threat detection and response both ineffective and inefficient.

Accuracy

The overwhelming tasks of verifying the quality and relevance of intelligence makes actionable threat intelligence, at mass scale, unattainable in practice. This results in a large amount of false-positives, leaving security analysts and responders chasing dead leads. EDR tools also leverage much of the same methods for detecting threats as defensive solutions such as NextGen AV and Endpoint Protection Platforms (EPP). They all attempt to catch the threat on the way in or to detect it by looking at the event logs. The accuracy of such detection methods is limited as they operate in-line, forcing them to sample at approximately 30-150 samples per minute. These approaches all produce not only false-positives but also false-negatives., meaning they are limited in what they can detect. pieces of logs, making the log-based approach to threat detection and response both ineffective and inefficient.

Resource Overhead

A security analyst that spends any of their time chasing dead leads, means an inefficient use of resources. Collecting intelligence and retaining logs also requires an overwhelming amount of system resources. Using security analytics then requires the use of a large data set and parallel processing to detect anomalies which may or may not be actual threats. Such practices lead to unoptimised risk and resources and leave businesses operating with limited detection capabilities.

Dependance

EDR and SI approach to threat detection, depend on matching pre-defined behaviours, tactics, and signatures. They attempt to detect the actual threat itself on the way in or rely on secondhand information as apposed to trying to detect the breach by looking directly at the actual evidence on the system itself. They do not take an independent approach to threat detection because they still rely on the same methods used by defensive solutions.

Forensic Depth Analysis Defined

Forensic-Depth Analysis (FDA) is an automated approach to post-breach detection that assumes devices are already compromised and seeks to validate every endpoint as thoroughly as possible. The automation inherent in FDA enables advanced threats to be detected effectively and at scale.

FDA operates independently from the host OS and uses dissolvable endpoint surveys to quickly collect live forensic data from both volatile and nonvolatile memory. Non memory-based information is also collected to identify persistence mechanisms. This data is then analysed using a variety of post-breach analytic techniques, and then enriched using multiple threat intelligence and reputation sources. Combining this live host forensic data and these analytic techniques, FDA determines the compromise depth of endpoints.

Defence in Depth Has Proven Gaps

Forensic State Analysis Targets What Defensive Technologies Miss

Anti-virus
Security Intelligence (SI)
Security Information & Event Menagement (SIEM)
User Behavior Analytics (UBA)
Endpoint Detection & Response (EDR)

Forensic-Depth Analysis (FDA) is a proactive solution used to hunt malware and persistent threats, whether these are on disk or fileless, and, does not rely on data from defences that have already failed. CyberStash, using FDA, surveys endpoints to determine compromise status and takes response actions to illuminate the threat before they lead to business impact.

Forensic-Depth Detection

CyberStash employs a Forensic-Depth methodology which is the most effective solution to determine the compromise depth of endpoints. CyberStash platform uses agents or dissolvable agents to independently collect, identify and evaluate a variety of data points, then analyses the data using forensic analytics and file intelligence services. Here are some of the functions CyberStash engages in:

Evaluating

- All active processes, loaded modules, scripts and drivers
- All Active Host Connections (including inter-process and redirects)

Identifying Disabled Security Controls

- Disabled AV
- Reduced authentication requirements
- GPO blocking, etc.

Identifying and Evaluating

- Memory Injected Modules – CyberStash FDA uses memory un-mapping techniques to export memory objects for offline retention and analysis
- Process Manipulation (such as function Hooks, inline modifications/patching, etc.)
- Operating System Manipulation (including list modifications, hidden processes, direct kernel object manipulation)

Enumerating and Evaluating Persistence Mechanisms

- Cronjobs
- Registry autostarts/triggers
- DLL hijacking
- WMI Events
- Boot process redirection
- Watchdog processes, etc.

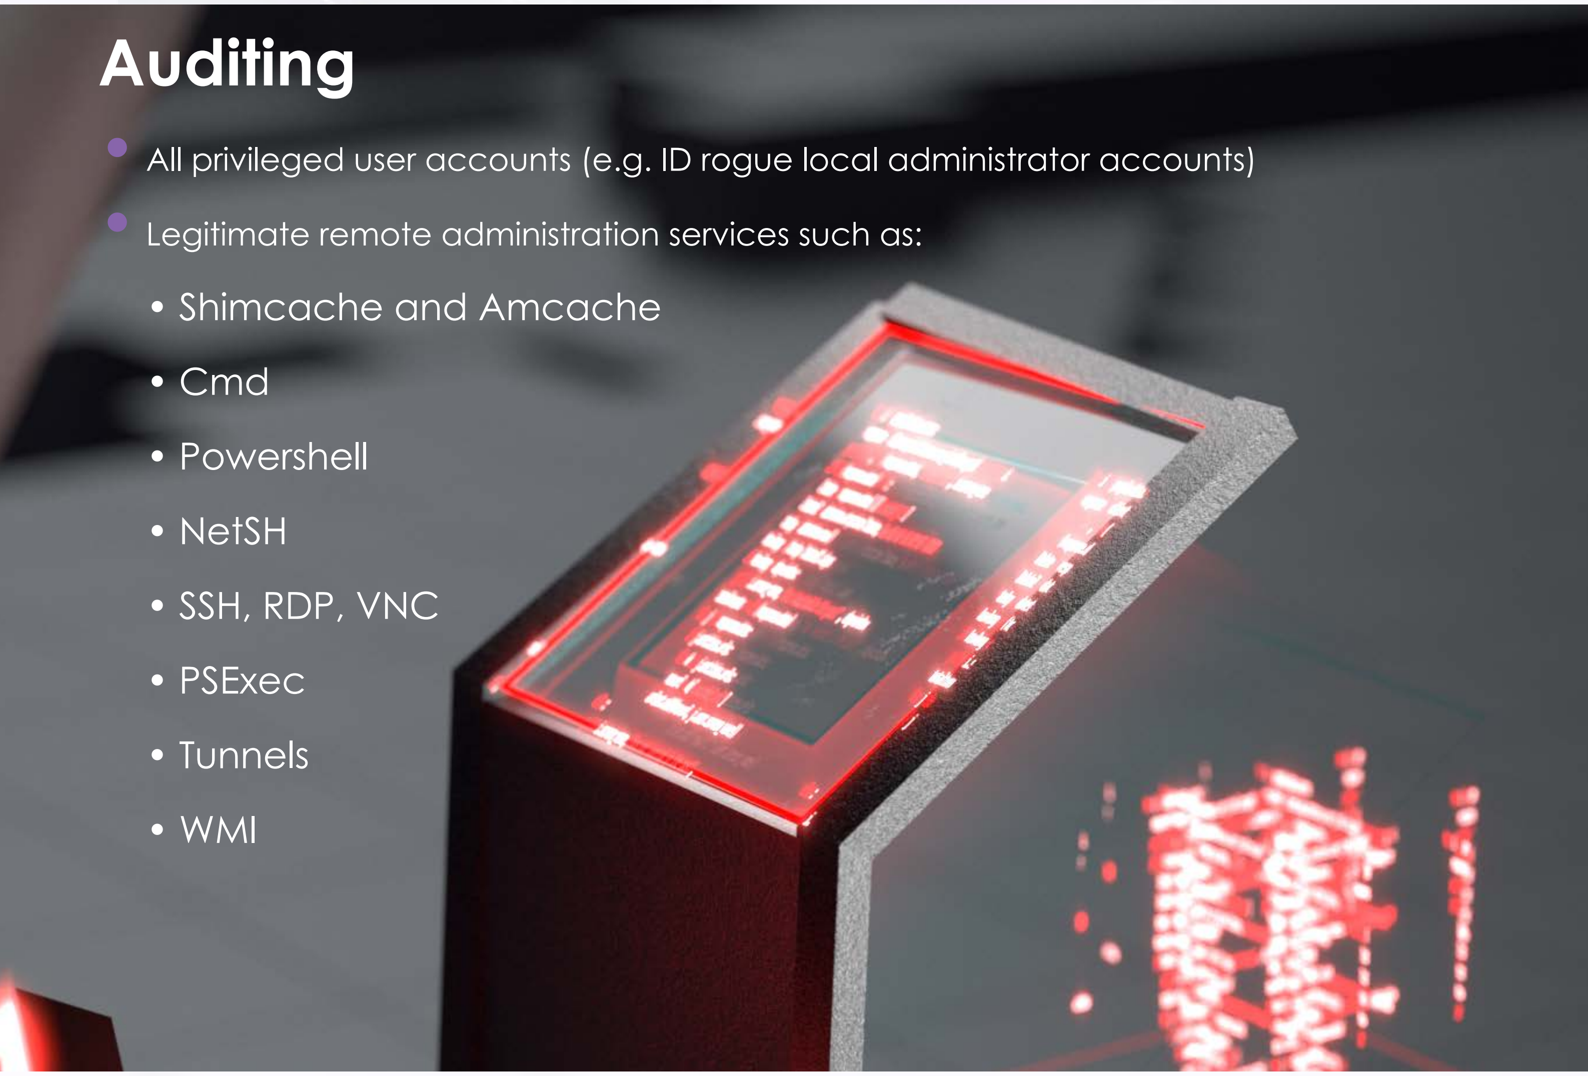

Auditing

- All privileged user accounts (e.g. ID rogue local administrator accounts)
- Legitimate remote administration services such as:
  - Shimcache and Amcache
  - Cmd
  - Powershell
  - NetSH
  - SSH, RDP, VNC
  - PSExec
  - Tunnels
  - WMI

Controlling Dwell Time

Business Impact

Without an effective advanced threat detection and response capability, organisations leave open the likelihood of business impact to occur. Organisations can reduce the likelihood of business impact by controlling dwell-time.

What Is Dwell-Time?

Dwell-time is the duration of time between when an organisation is breached by a cyber-attack to when the breach is discovered and removed. With dwell-time averaging from about 6 months in EMEA to 7 months in Asia Pacific, limiting dwell-time is a key performance indicator that must be controlled. Controlling dwell-time to 21 days, reduces business impact by 40%, with exponentially increasing benefits the further dwell-time is reduced. Controlling dwell-time to 1 day, delivers a 96% reduction in business impact.

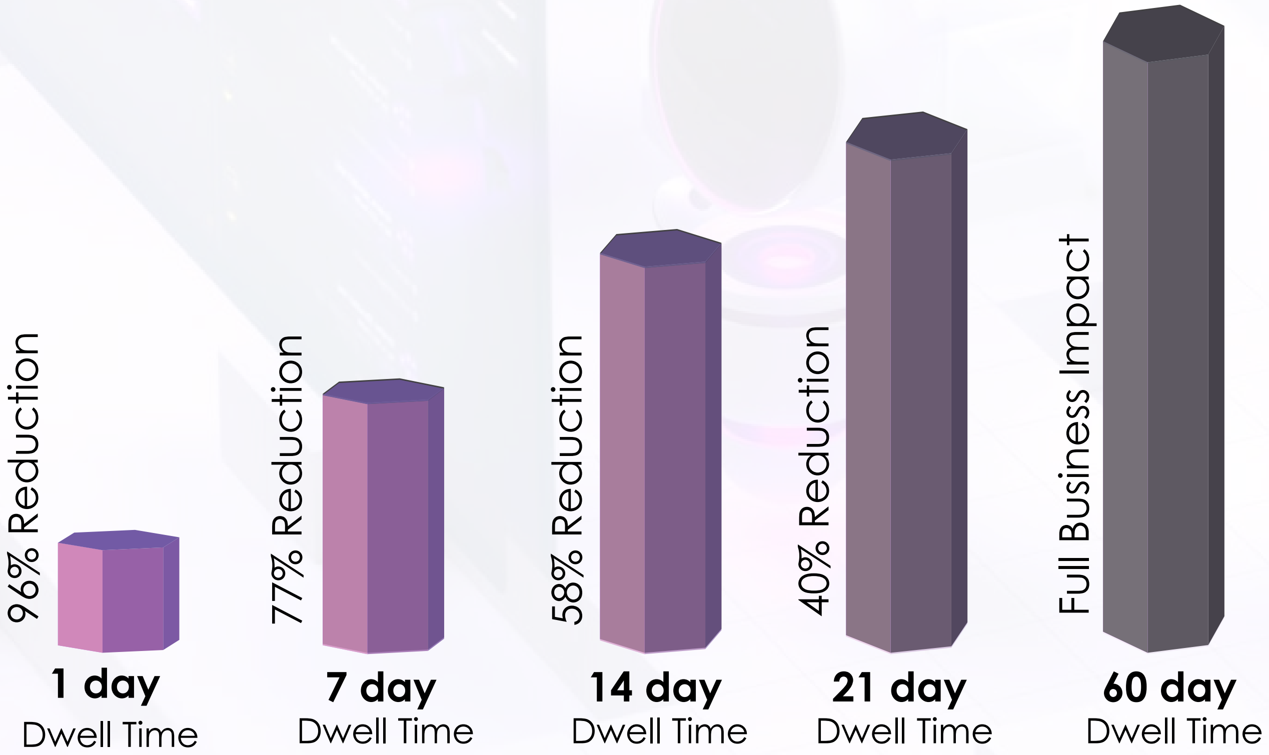

How To Reduce Dwell-Time And Avoid Business Impact Without Adding Resources

Using Forensic-Depth Analysis (FDA), CyberStash systematically surveys your endpoints to determine whether your organisation has been breached. Our automated approach to Forensic-Depth Analysis (FDA), delivers a post-breach defence strategy to businesses to increase their resilience to cyber-attacks. Endpoints found to be compromised are then flagged for isolation or clean-up.

This is made possible by the sheer depth of analysis which includes an advanced survey of a host's volatile memory, application persistence mechanisms, forensic artifacts and a thorough verification of operating system (OS) integrity. CyberStash takes the art of memory forensics to a new level of scalability by surveying the live memory of thousands of endpoints, simultaneously.

Use Cases Why Choose Cyberstash

The CyberStash Post-Breach Forensic-Depth Compromise Assessment Service, provides an independent audit of your IT assets to ascertain whether any business systems are currently breached. 'Independent', not only because we are an external 3rd-party, who is carrying out the assessment, but because of the methods we use to conduct the assessment is completely independent of the existing toolsets used within your environment to detect threats. Namely, we use Forensic-Depth Analysis (FDA).

The use cases for conducting Compromise Assessments are:

- Independently verify the current security posture of your IT environment.
- Detect and respond to advanced cyber breaches that have circumvented your existing controls.
- Following a cyber incident, validate that all human adversaries, backdoors and malware have been completely cleaned out.
- Build resilience by controlling dwell-time, reducing risk and maintaining compliance.
- Provide cyber assurance to business stakeholders to reindepth their trust and confidence in IT systems.

Want to run a trial?

Reach Out To Cyberstash
For More information.

Explore

eclipse.xdr

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