Reducing Attacker Dwell Time

How to stop attackers before they can destroy or steal your critical information and IT assets

WHITE PAPER

Infocyte®

Table of Contents

Executive Summary

Our networks are attacked hundreds, sometimes thousands, of times a day by hackers and fraudsters throughout the world. Occasionally, these attacks are successful in gaining a foothold into the targeted network. Worse, skilled attackers have demonstrated they can remain hidden for months, sometimes years, before detection once inside. When able to maintain long term, persistent access attackers can spy on operations; steal sensitive information; corrupt files; and even cause physical damage by manipulating industrial control systems (i.e. motors, actuators, or power).

According to the Ponemon Institute the time to identify and time to contain a breach is highest for malicious and criminal attacks (229 and 82 days, respectively).[1]

According to several industry reports, the average network security breach goes undetected for 6 months on average. According to the Ponemon Institute's 2016 global study, the time to identify and time to contain a breach is highest for malicious and criminal attacks (229 and 82 days, respectively).[1] This problem is known in the security industry as "attacker dwell time" and it represents one of the greatest threats for any organization that uses information technology.

Today's detection methods and technologies have predominantly focused on the real-time prevention and detection of attacks through 24/7 monitoring. What is missing are processes and technology that address detection of adversaries and insider threats that are already in the network or on a connected device. This type of post-compromise detection has typically been the realm of reactive incident response, but those that want to be proactive have another option; "threat hunting".

This white paper introduces threat hunting as an essential component of your defense-in-depth strategy to combat persistent compromises and hidden cyber threats. We will examine why adversaries are successful in persisting in networks; the limitations of existing security technologies and methodologies to discover threats once they are inside; and how dedicated hunt technology and processes can work with your existing security infrastructure to deny attackers the ability to persist undetected.

[1] 2016 Cost of Data Breach Study: Global Analysis, Ponemon Institute

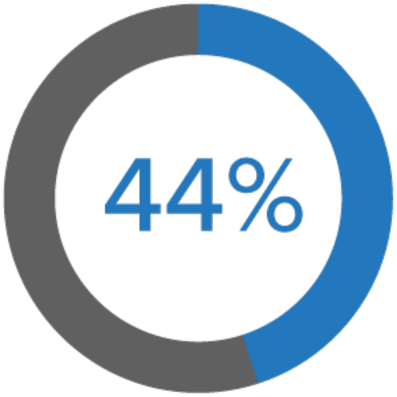

Threats undetected by automated security tools.[2]

Understanding Attacker Dwell Time

Recent attacks reported by the media continue to highlight various breaches that have gone undetected for weeks, months and sometimes years. Known as attacker dwell time, this period is defined as the time elapsed between the initial breach of a network by an attacker and the discovery of that breach by the victim.

This trend does not show signs of slowing as internal security processes and tools are unable to keep up with an increasingly sophisticated and pervasive threat. According to the 2017 Threat Hunting Report, 44% of threats go undetected by automated security tools.[2] And a Trustwave Report showed that 81% of reported intrusions are not detected by internal security processes, but rather by news reports, law enforcement notifications, or external fraud monitoring.[3]

The examples in Table 1 offer a snapshot of recent real-world attacks and the length of time elapsed before the breach was discovered. These well documented incidents cost the organizations affected millions in losses, regulatory fines and brand reputation.

Table 1: Real-World Attacker Dwell Time Examples

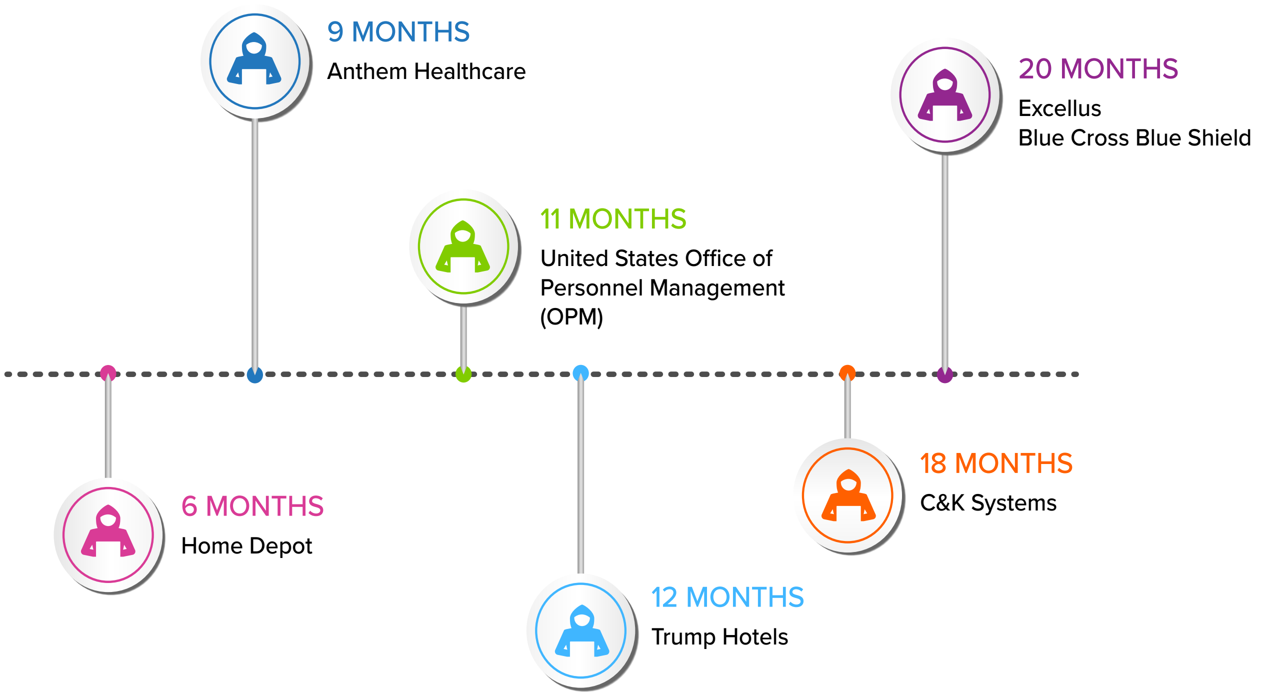

[2] 2017 Threat Hunting Report, Crowd Research Partners

[3] Trustwave Report as reported in ComputerWorld UK, http://www.computerworlduk.com/news/security/most-data-breaches-still-discovered-by-third-parties-3615783

Known as “persistent compromises" there are many motives for attackers trying to maintain stealthy long term access to a network. Whereas loud, transient attacks like crypto-locker, web defacement, denial of service, or smash and grabs can be easy to identify due to the immediate effect they have, persistent threats meet their objectives by maintaining stealthy long term access to the network.

Table 2: Persistent vs Non-persistent Compromises

Persistent Compromises	Transient Compromises
Spying	Extortion (i.e. Crypto-Locker)
Corporate Espionage	Web Defacement
Credit Card Theft	“Smash and Grab” Theft
Botnet Operations	Denial of Service
False Flag Attacks & Pivots	Destructive Worms
Posturing for Future Attack (i.e. military)	

“IT organizations lack the ability to detect issues and spot early warning signs that malware has slipped past preventive measures.”

- Gartner

While access may be obtained within seconds or minutes depending on the vulnerability exploited, mapping and navigating a large or complicated network to find the data or individuals the attacker is looking for can many times take days or weeks. Additionally, monitoring users on the newly compromised network for a period of time to learn internal operations is essential to an attacker’s success, as was demonstrated in the Sony attack. This however also gives network defenders an opportunity to disrupt and counter.

Strategies to Reduce Dwell Time

Although the attacker dwell time problem is complex, it exists primarily for two reasons:

1. The growing sophistication of modern attackers.
2. Current real-time security processes are ineffective at detecting post-compromise activity, especially as time passes after the initial attack.

The issue of hidden, persistent compromises has become so pervasive that many argue organizations should operate under the assumption that their respective networks will be penetrated, if they aren’t already. The U.S. Department of Defense adopted this premise several years ago, and in response, created "hunt teams". At a basic level the teams consist of trained incident responders and analysts who proactively and iteratively search critical networks and/or historical log data for signs of a missed compromise.

[4] Lee, Robert M., Lee, Rob, The Who, What, Where, When, Why and How of Effective Threat Hunting, SANS, Feb 2016, http://www.sans.org/reading-room/whitepapers/analyst/who-what-where-when-effective-threat-hunting-36785

Threat Hunting

While first developed by the military, enterprises are now begining to see the value in this approach and adopting threat hunting practices to root out attackers. Threat hunting is defined by the SANS Institute's Rob Lee as:

> "A focused and iterative approach to searching out, identifying and understanding adversaries internal to the defender's networks"[4].

Threat hunting is differentiated from real-time intrusion detection, which works to prevent or detect attacks early in the attack cycle, by instead utilizing post-compromise detection techniques. Hunting is on the spectrum of incident response activities except it is done proactively, before you know there is a problem. The goal is to reduce the dwell time of attackers and remove them before they can cause further damage.

You can keep hardening the front door, but that does nothing against the threat already inside.

A couple approaches to threat hunting that have been adopted by the security industry are:

1. **Anomaly detection and analytics** on existing logs, connection, and event data to catch what real-time analysts might have missed.
2. **Monitor the network with new/different security sensors** (i.e. passive DNS monitoring).
3. **Actively scan devices for indications of compromise using an endpoint-focused solution** (i.e. IoCs, multiple AV engines, artifacts & malicious behavior).

Each of these approaches have pros and cons, with varying costs and effectiveness. Analytics can be effective but require existing high quality security data that goes back far enough (i.e. if the initial breach happened one year ago, 30 day log retention may not assist you in finding that). Monitoring using additional tools can be cost prohibitive and may give overlapping coverage reducing its' overall value and effectiveness.

We'll go more into the pros and cons of each as well as the approach we recommend, but first we need to look at why traditional security monitoring isn't doing this job.

Traditional Intrusion Detection

The three pillars of a traditional network security program consist of preventing, detecting, and responding to intrusions on a network.

- **Prevention** is performed through patching, IP and domain filtering, application whitelisting, and various network or host based intrusion prevention systems.
- **Detection**, thus far, has focused on continuous real-time monitoring conducted using various network and host-based sensors. These sensors can identify attacks based on previously seen characteristics (signatures) and behaviors (heuristics). Their primary focus is to detect attacks as early as possible to limit the damage from a successful intrusion. Detection generally takes place in two broad areas:
    - *Network-Based* – Use of network traffic inspection sensors at network choke points and internet gateways. Examples include firewalls, intrusion detection systems (IDS), and proxies.

    - *Host-Based* – Inspection of endpoint device activity and data. Examples include Anti-virus and Host Intrusion Detections Systems (HIDS).

- **Response** has traditionally centered on the post-mortem forensic investigation performed by an incident response professional or team. Standard actions when an intrusion does occur involve the characterization of the attack via forensics, implementation of network blocks to prevent further activity, and restoration of compromised systems to a clean state. The forensics process is typically time consuming and expensive for organizations, but it uncovers much more detail than is available from real-time detection sensors.

Unfortunately, present day attackers are proving more agile than the above-mentioned processes. Detection capabilities have been forced to progress beyond reliable signatures into noisy heuristics to detect increasingly stealthy adversaries. Meanwhile, incident response continues to be prohibitively expensive for many organizations. This has made incident response and forensics an emergency resource for many organizations, as only a confirmed compromise with demonstrable impact to the business can justify the expense.

Threat Hunting has been marketed by some as a very intense process requiring high skillsets and detail oriented analysts. This is simply due to a lack of automation and solutions tailored to the activity.

Real-Time vs Post-Compromise Detection

Another issue is that the tools and techniques required to detect successful compromises and adversaries already inside the network differ from those required to detect attacks in real-time.

Real-time tools and strategies focus on early detection of attacks, exploits, malware installations events to prevent an attack from succeeding or catching it early enough to mitigate the damage. While important, most indicators associated with the real-time paradigm are present only during the initial attack. If an attacker successfully evades real-time protection, the attacker will become entrenched by installing rootkits, establishing encrypted /alternate command and control (C2) channels, and stealing credentials to better blend in with the network, rendering the majority of real-time intrusion detection techniques ineffective.

Threat hunting answers the need to identify ongoing or successful compromises that have made it past these traditional defenses undetected. The tools and techniques used in threat hunting must be able to identify post-compromise activity, dormant or hidden malware, malicious use of credentials, and Command and Control (C2) traffic. While this practice of post-compromise detection is on the spectrum of incident response (IR) activities, traditional IR tools are too limited by complexity and lack the scalability to perform proactive hunt. Some of the drawbacks of existing IR and forensic tools include:

- Overly time consuming and manpower intensive.
- Not scalable past a single or handful of systems.
- Multiple disparate tools are required.

Although it is possible to hunt for threats by combining various IR tools, the need to reduce time and cost demands a purpose-built solution dedicated to post-compromise threats.

Threat Hunting has been marketed by some as a very intense process requiring high skillsets and detail oriented analysts. This is simply due to a lack of automation and solutions tailored to the activity. An automated hunt workflow can multiply the effectiveness of an analyst by orders of magnitude allowing a small team to cover an entire enterprise network with hundreds of thousands of nodes. Additionally, integrating existing detection and analysis technologies like threat intelligence databases and malware sandboxes will be beneficial.

Infocyte approaches threat hunting through the independent scanning and validation of endpoint devices (workstations, servers, mobile, etc.). Because persistent threats must maintain access or go through an endpoint device in order to navigate the network, it is here that we will find the most effectiveness in discovering unauthorized access, malware, and indications of compromise.

Network-based hunt is limited by the fact that attackers use encrypted communications and skilled ones will attempt to blend in to the noise of traffic on the network making anomaly detection ineffective. Log and event-based analytics solutions suffer due relying on the existing security infrastructure to provide the right data. Most organizations do not even retain logs long enough to identify breaches that occurred months or years ago. Finally, many analytics solutions still suffer from low signal to noise ratios.

While a combination of the techniques can be effective, it ultimately has diminishing returns. We have found through years of hunting that endpoint scans discover everything we could find from other solutions and more, making it the most cost effective approach to finding existing compromises. For this reason, we built Infocyte HUNT as an agentless endpoint scanner.

Introducing Infocyte HUNT™

Infocyte is used by an organization's own hunt/security teams and incident responders to find hidden threats and reduce the dwell time of hackers who have made it inside the network. Being agentless, we don't monitor endpoints the same way your endpoint protection suite works, instead, we utilize a set of highly scalable Digital Forensics and Incident Response (DFIR) methods called Forensic State Analysis (FSA). Using FSA, our Infocyte HUNT platform periodically sweeps thousands of endpoints, spending a couple minutes on each host, to definitively validate their state as either "Compromised" or "Not Compromised" with greater confidence than antivirus or intrusion monitoring can provide.

Infocyte HUNT digs deep into an endpoint to validate:

- What is actively running in memory;
- What is triggered to run through a persistence mechanism, and;
- Identify any manipulation that would suggest the system has been maliciously modified (e.g. what a rootkit does to hide its presence, or what an insider threat might do to disable the system's security controls).

Conclusion

The majority of detection capabilities today focus on real-time attack detection and prevention, and incident response continues to be out of reach of most organizations. What's needed is a layered approach to breach discovery that includes looking for threats which have slipped past front line defenses.

Organizations need to employ threat hunting as part of their security strategy in order to catch what prevention and real-time detection technologies miss to mitigate the possible damage that can be caused from prolonged unauthorized access. Threat hunting moves network security operations away from a rigid sensor and investigation-centered process to a more proactive defense model.

The Infocyte HUNT solution enables any company's security team to eliminate attacker dwell time by making it easy, accurate and cost-effective to find and remove attacker presence long before they can achieve their objective.

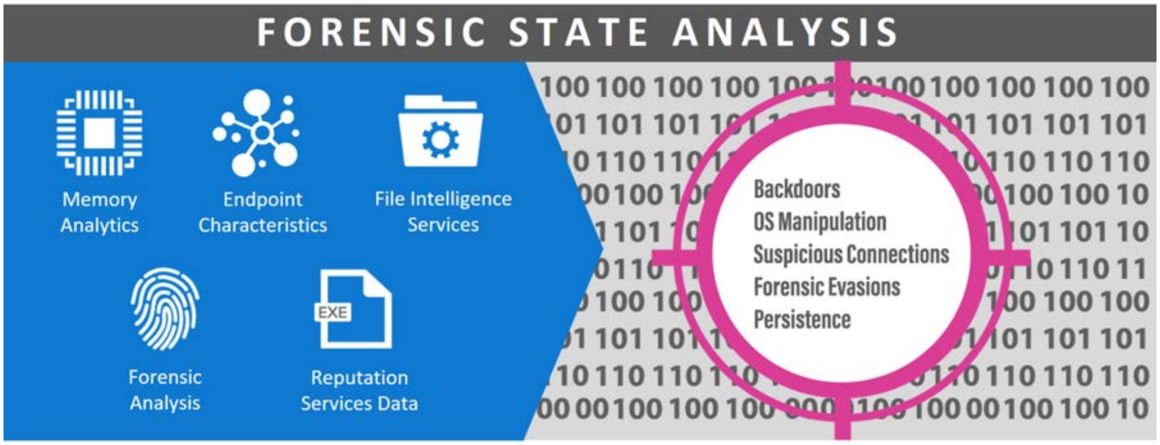

For more information or to request a demo of Infocyte HUNT go to: www.infocyte.com

Infocyte®

About Infocyte

Developed by former US Air Force cybersecurity officers, Infocyte's hunt technology fills a void left by today's real-time detection solutions. By focusing on the post-compromise activity of persistent attackers and insider threats, Infocyte's unique approach to security helps organizations defend their networks and critical information.

CORPORATE HEADQUARTERS
110 E. Houston St. Floor 6
San Antonio, TX 78205
+ 1.844.INFOCYTE (844.463.6298)
sales@infocyte.com

www.infocyte.com
@InfocyteInc