VECTRA
SECURITY THAT THINKS

WHITE PAPER

How the Cognito platform replaces IDS and enables organizations to detect intrusions again

TABLE OF CONTENTS

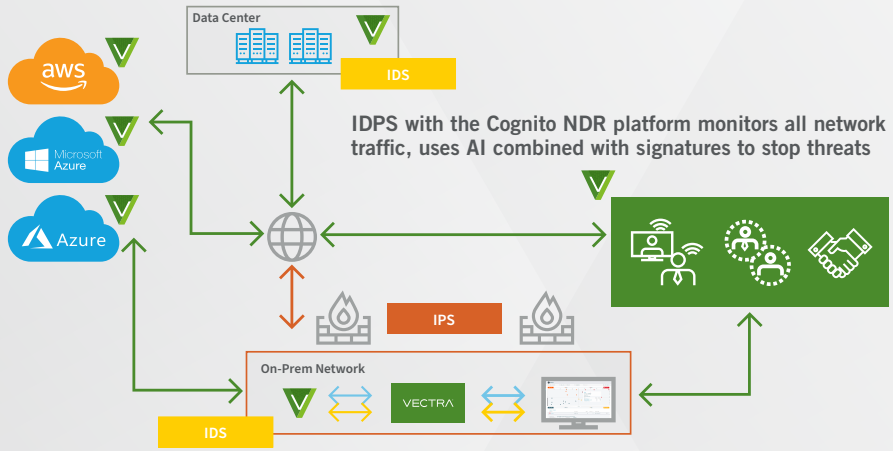

HIGHLIGHTS

- Attackers have adapted to the enterprise security model and can easily evade and avoid perimeter and malware detection techniques.
- Cognito applies algorithmic models directly to network traffic to reveal underlying attack behaviors and then enriches that data with secondary sources such as authentication logs and threat intelligence data – automatically.
- Cognito uncovers in-progress cyberattacks inside a network. Detections are correlated to the hosts under attack, and each is scored and prioritized according to the highest risk.

Introduction

Intrusion detection systems (IDS) have been a mainstay of information security for decades.

However, the "Gartner Market Guide for Intrusion Detection and Prevention Systems" by Craig Lawson and John Watts, published July 1, 2019 (ID: G00385800), notes that, "The plethora of breaches continues unabated, which highlights how organizations need to better address the protection of internal assets and improve their ability to detect and prevent the lateral movement of threats."

Today, standalone IDS has been subsumed by intrusion prevention systems (IPS), and the two are now known collectively as IDPS. This convergence occurred as the security industry focused more on preventing external threat actors, largely due to the lack of skilled security analysts able to make sense of the volumes of noise presented by IDS.

One must assume a breach has already occurred and detect what perimeter systems missed, including IPS.

However, this mode of thinking is dangerous, as evidenced by the huge increase in large breaches over the last few years. IPS is externally focused as an extension of the perimeter firewall, where the goal is to stop malware and external threats from entering the network.

As persistent cyberattacks and network breaches become more common, the need for intrusion detection is higher than ever before. Unfortunately, existing IDS techniques fail to meet today's security requirements.

Detection of attacker activity in the cloud and internal to the network is not the same as detection of the initial infection, which is usually malware. In fact, one must assume a breach has already occurred and detect what perimeter systems missed, including IPS.

"The plethora of breaches continues unabated, which highlights how organizations need to better address the protection of internal assets and improve their ability to detect and prevent the lateral movement of threats."

Craig Lawson and John Watts
July 1, 2019 Gartner Market Guide for Intrusion Detection and Prevention Systems" (ID: G00385800)

This white paper the Cognito® Network Detection and Response platform from Vectra® – which automates the hunt for cyberattackers in cloud, data center, IT, and IoT networks – is ideal for replacing aging IDS. It starts with a brief history of intrusion detection and how it led to the limitations IDS faces today.

The white paper also explains how to move intrusion detection forward, including the unique challenges of detecting modern threats and the new detection technologies and architectures designed to solve the problem.

A brief history of IDS

Intrusion detection dates back to the early 1980s and the pioneering work of Dorothy Denning and Peter Neumann. Research into IDS was driven largely by the U.S. government, which sought to protect confidential assets from internal users. This is a vital distinction because threats were defined more by misbehaving or internal users and not external attackers. At that time, all attacks were essentially insider attacks.

Many concepts behind the first IDS remain relevant today. The goal was to build rules that reveal suspicious behavior and identify deviations from normal baselines. They relied heavily on establishing baselines and finding anomalies by analyzing audit logs at the host level.

Today, IDS is deployed on the host or the network, with host-based IDS monitoring a single host and network-based IDS monitoring the entire network. The two dominant detection models are signature-based and anomaly-based.

First introduced in antivirus technology, signatures detect attacks by looking for specific patterns, such as distinct byte sequences in network traffic or malicious code used by malware.

Anomaly models detect attacks by looking at what is different on the network and hosts, such as user authentication from unrecognized systems or large volumes of incoming data from previously unknown sources.

Signature and anomaly models are widely used today, despite fundamental weaknesses. Neither model addresses longstanding key limitations of IDS, which requires highly skilled analysts to interpret large volumes of data in a time-sensitive manner. This alone has adversely impacted the real-world effectiveness of IDS.

The shift from IDS to IPS

Over time, IDS had an architectural shift in its use and deployment to address threats from outside attackers. Despite this shift, IDS remained time consuming and resource intensive.

Parsing through vast amounts of data from IDS required highly-skilled security analysts to find the real attacks in the midst of all the network noise. With time being a critical factor, the longer a threat was able to progress undetected, the more time an attacker had to steal critical assets and cause damage. This remains true today.

As the internet grew, more threats came from outside the walls of the network, and security teams needed a way to defend their applications and assets. As security teams shifted their focus to detecting and stopping external threats, IDS vendors followed suit and found better market success.

IDS was deployed at the perimeter, after the firewall, with no further internal visibility. At the time, finding attackers as they enter and leave the network was thought to be easier than hunting for hidden attackers already inside the network. Despite its lack of internal network visibility, limiting the scope of detection to specific external threat actors reduced the security analyst workload.

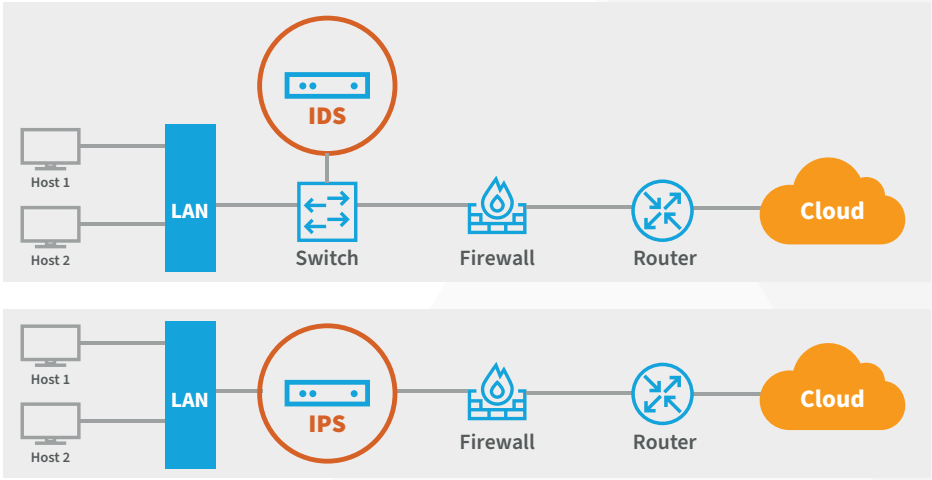

As the Internet grew, more threats were coming from outside the walls of the network.

Although perimeter firewalls filter network traffic to ensure only authorized applications and users have access to certain resources, they do not have visibility into malware, threat behaviors and credential-based attacks.

As security teams shifted their IDS budgets to external threat detection at the perimeter, IDS vendors shifted their focus to this specific use-case and enabled inline blocking on a subset of detections.

This became known as "lean-back security" with a reduced workload on the security analyst. IPS was thus born, and rapidly became a critical and standard layer of the network perimeter, where it evolved into a function of today's next-generation firewalls.

IPS and the hidden impact on intrusion detection

There are valid reasons for focusing on IPS. The internet harbors a virtually unlimited number of threats and security teams need a way to block the vast majority of them automatically. If these threats can be blocked before they enter the network, IPS can theoretically reduce the need for labor-intensive threat hunting.

Speed is king

Speed is critical for perimeter-based IPS. It cannot sacrifice network availability as it performs security monitoring. It must never block approved users from accessing networked resources. High speed, fast throughput, and low latency are non-negotiable. These demands have led to constant trade-offs that sacrifice detection capabilities for speed.

Detection must be extremely fast for prevention to be feasible. When decisions must be made in milliseconds, there simply isn't time to think. Detection and the corresponding response must be near-instantaneous and reflexive. Detection must also be highly accurate to be enabled for blocking.

This has created problems for signature-based detection models. While signatures come in many forms – exploit focused, vulnerability focused, malware hashes, known bad IP addresses, and known bad URLs – they all depend on very fast pattern-matching of known threats.

To achieve its goals, signatures require a trade-off between the volume of detections that can be enabled with the required throughput performance. As detections increase, performance declines, which forces security teams to pick and choose signatures to strike the right balance of security and speed.

This means IPS only provides protection against the most well-known attacks. It also creates a dangerous security gap between the time a threat is discovered in the wild and the time IPS can confidently respond.

Putting detection first

IPS and IDS have very different use cases and roles in the network. While IPS is tuned for performance, it is limited to a subset of detection techniques that can be performed quickly and with limited memory.

In today's threat landscape, where high-profile breaches are frequently reported in the news, it is clear that prevention techniques alone are insufficient. Detecting and responding to hidden attacks that progress inside cloud, data center, IT, and IoT networks must be a top priority.

To meet today's challenges, NDR needs to understand the way attacks really work and how attackers succeed. Today's sophisticated attackers are armed with the same tools used by system administrators and do not have to use malware or exploits.

NDR must adapt and reflect the true nature of cloud, data center, IT, and IoT networks and their attack surfaces. Endpoints are dynamic and increasingly mobile, servers can be hosted inside the network or in the cloud, and security analysts have an increasingly difficult time with asset management and knowing where data resides.

Today's sophisticated attackers use customized encryption that can't be decrypted, even under the best of circumstances.

Most importantly, NDR cannot be complex. It must be easy to deploy, manage and use. It should not require a full-time expert to keep it operational. To ensure real attacks don't go unnoticed in large volumes of noise, NDR must implement a method of noise reduction and risk prioritization that enables a security analyst to focus on the threats that pose the biggest risk.

Modern threats and the need for true intrusion detection

Attackers today can easily evade and avoid perimeter and malware detection techniques. Evasion involves getting threats past the perimeter without detection, while avoidance finds avenues into the network without crossing boundaries.

Signature evasion

The most straightforward approach to evading signature-based IDS is to use traffic that doesn't match known signatures. Depending on the signature, this can be trivial or highly complicated. For example, signatures based on known bad IP addresses and URLs are often used to identify command-and-control servers of botnets and malware. For attackers, avoiding signatures is as easy as registering a new domain.

At the other end of the spectrum, highly-sophisticated attackers can find and exploit previously unknown vulnerabilities. Attacks on such unknown, or zero-day, vulnerabilities naturally lack signatures because they are unknown to the security industry.

Attackers can scramble the attack payload, making it difficult for IDS to recognize it.

Other signature evasions confuse the signature match in a variety of ways. Attackers can scramble the attack payload, making it difficult for IDS to recognize. Fragmenting and reordering are widely-known techniques that IDS/IPS systems are prepared to catch, but there are near-infinite numbers of evasion combinations and tricks that attackers can use to sneak through.

Encryption

Another way to avoid signatures is to obscure the traffic. Instead of developing an exotic new exploit, it's easier for an attacker to ensure that security doesn't get a good look at the traffic. This can be as simple as encrypting malicious network traffic.

To complement the growing use of SSL/TLS encryption, attackers use a variety of applications in their arsenal that are encrypted by default to ferry malicious traffic past the perimeter. While SSL decryption at the perimeter is an option, it's costly, introduces performance penalties, and has become complicated due to industry changes like certificate pinning.

Today's sophisticated attackers use customized encryption that can't be decrypted, even under the best of circumstances. This leaves security teams in the unenviable position of either blocking or allowing unknown traffic at the perimeter.

Perimeter avoidance

In addition to techniques that smuggle cyberthreats across the perimeter, attackers have learned to avoid the perimeter altogether. By infecting users' devices at home or outside the perimeter, threats can be carried in through the front door.

Mobile devices provide logical and physical paths around the perimeter. Mobile devices with LTE or 5G data connectivity have easy paths to the internet and can introduce serious risks that do not cross the network perimeter. It's an invisible conduit that attackers love to use to get inside networks.

Inside moves

Given the almost exclusive focus of IDS/IPS on the perimeter, attackers can move much more freely once they are inside. This allows attackers to develop patient and methodical attacks that gradually extend through a network in search of key assets.

This involves an ongoing process of internal reconnaissance, lateral movement, and the access and theft of key assets. Each area involves a wide variety of techniques and strategies on the part of attackers, and they all take place inside the network where visibility is typically low.

As organizations move their high-value data and services to the cloud, it's imperative to reduce cyber-risks that can take down businesses. Visibility gaps can exist in connections between compute and storage instances.

Cyberattackers are aware of this visibility gap. A recent survey by the SANS Institute found that one in five businesses had serious unauthorized access to their cloud environments this past year alone, and many more were unknowingly breached. This will only become more pronounced as nearly four out of 10 organizations plan to move to a cloud-first approach to deploy new applications, according to a recent study by the Enterprise Strategy Group (ESG).

Moving beyond exploits and malware

Once inside the network, savvy attackers don't need exploits and malware to extend their incursion. Instead, they simply harvest user credentials from compromised hosts to spread through the network.

This can be done by capturing a username and login during the authentication process or stealing credentials or hashes from memory. In either case, attackers can spread throughout the network using valid credentials without having to use exploits or malware.

Attackers can also blend in as trusted users and leverage any number of applications. Webmail, social media, and virtually any type of browser or web-based application provide a conduit between attackers in the network and the outside world.

After quickly learning which applications and tools are used in the network, attackers will add them to their arsenal. For example, remote desktop applications or file-sharing applications like Dropbox can be powerful attacker tools that are allowed by network policy.

NDR is defining the next generation of security

Modern requirements for NDR will expand to include points of visibility in the cloud and internal network but also identify surgical behaviors rather than generate volumes of alerts.

The importance of cloud and internal network visibility

Attacks can be classified as one of three types: Targeted, insider or opportunistic. Knowing one from the other requires a full understanding and context of what the attacker is doing inside the network, especially as malware is shared by attackers and usually only represents the initial infection component.

After the initial infection occurs, a security analyst is required to make a very fast critical decision about how to respond.

IDS has little hope of detecting threats if it can't see where the majority of action is taking place.

As a result, it's more important for NDR to monitor what is happening inside of cloud footprints than what has entered the network. Internal reconnaissance, lateral movement, unauthorized data access, and staging will occur across any part of the infrastructure inside the network between compute and storage instances, user devices and servers.

Without full traffic visibility, security teams are limited in their ability to see the entire attack lifecycle, which in turn limits the understanding and context of what is really happening.

Instead of deploying IDS/IPS at ingress and egress points, NDR should be deployed to provide the greatest visibility to watch all user-to-user, user-to-server, server-to-server, and user-to-internet traffic. The goal is to monitor all traffic and behaviors across every asset in cloud, data center, IT, and IoT networks.

Finally, NDR should provide context and prioritize the behaviors it sees, including options for responses relevant to specific threats. Because every attack is different, every response should be different. This is best addressed by integrating and automating best-of-breed threat-response and containment technologies already in the network, such as endpoint protection, firewalls and network access controls.

New models of attacker detection

Moving from payloads to behaviors

To detect sophisticated attacks, NDR must move beyond the realm of solely relying on signatures and simple anomalous behavior detection. An underlying limitation of signatures is that they typically search for malicious payloads, while anomaly detection only knows what is different instead of what is bad.

It's easy for attackers to adapt and avoid these types of controls. New exploits, modified tools, repackaged malware, and new IP addresses and URLs enable attackers to avoid detection. Attackers will also adopt normal user behaviors to avoid standing out from normal traffic and activity.

To move beyond this challenge, detection models should focus on identifying the underlying malicious behaviors, combined with curated indicators of compromise. This is conceptually akin to looking for malicious verbs as opposed to malicious nouns.

Detection models should focus on identifying the underlying malicious behavior.

This approach can be quite powerful. Although attackers can easily put on a new coat of paint to avoid signatures, they are exposed by their malicious behaviors. Attackers have a near-infinite supply of tools to help them to spy, spread, and steal inside the network, but they must perform the same tell-tale actions and behaviors to carry out an attack.

By learning to recognize the unique characteristics of these malicious behaviors, security teams can reliably identify network intrusions, even if the tools, malware or attack are completely unknown.

This level of detection requires a deeper understanding of malicious behaviors that goes well beyond the basic knowledge of an application. For example, an attacker may use a valid, allowed application such as RDP for command-and-control communication.

To detect threats, security technology must recognize the unique behaviors of command-and-control traffic, regardless of the application being used. The behavior of malware that requests instructions or updates the binary should be detected, whether it occurs via webmail, Twitter or Dropbox.

The science of detection

Powered by AI, the Cognito network detection and response platform from Vectra is based on the direct analysis of traffic to reveal the fundamental behaviors at the heart of cyberattacks.

By combining data science, machine learning and behavioral analysis with well-curated threat intelligence, Cognito identifies the intent of network traffic and reveals malicious behaviors, independent of applications and even when traffic is encrypted. This approach reveals the key actions that an attacker must perform to succeed.

Cognito applies algorithmic models directly to network traffic to reveal underlying attack behaviors and then enriches that data with secondary sources such as authentication logs and threat intelligence data – automatically. While these secondary sources are not required to detect an attacker, they provide context to accelerate the detection and response process for security analysts.

Cognito applies algorithmic models directly to network traffic to reveal underlying attack behaviors.

For example, an attacker using a custom remote access tool (RAT) can bypass traditional signatures and appear like a normal internet connection during an analysis of logs or NetFlow records.

By mathematically analyzing the connection at the packet level, Cognito identifies the unique pattern of an outsider who is controlling a machine inside the network. These concepts are explained in more detail in the next sections, which delve into the differences between supervised machine learning and unsupervised machine learning.

Different styles of machine learning

Detecting threats requires two types of high-level experiences. The first is a global set of experiences that understands how threats differ from normal or benign traffic. Second is a local set of experiences that reveals unusual or anomalous behaviors in a given environment.

The first approach reveals behaviors that are always bad and the second reveals threats based on local context. Both are essential to detecting threats, and they must work cooperatively.

Supervised machine learning addresses the former challenge by analyzing known malware, threats and attack techniques. Guided by Vectra data scientists who identify fundamental post-exploit behaviors that are consistent across all variants, this analysis feeds algorithms that detect underlying malicious behavior in network traffic.

While global intelligence is critical, some attacks are only revealed based on understanding the local context of the target network. Unsupervised machine learning refers to models that proactively recognize what is normal for a particular environment and when behaviors deviate from that norm.

Both styles of machine learning are essential and work together to detect hidden threats. Likewise, both styles support detection algorithms based on information that is observed over extended periods of time. Instead of detecting in a few milliseconds based on a single packet or flow of data, Cognito models learn and detect based on times ranging from seconds to weeks.

Supervised machine learning

When applied to network traffic, supervised machine learning gives security teams a big advantage. By applying it to large samples of post-exploit traffic and prevalent attack techniques, Vectra data scientists can identify the key traits they have in common. This enables Cognito to build algorithms that detect all variants of a style of threat.

By identifying common underlying behaviors, Cognito breaks the cycle that plagues signature-based solutions. Cognito knows the unique patterns of command-and-control servers that guide internally-infected hosts.

Supervised machine learning identifies fundamental behaviors that are consistent across all variants of a threat.

In the past, when new command-and-control signatures were released, attackers simply moved to a new server and proceeded without a problem. Cognito continues to detect new variants as well as completely new types of malware.

Vectra data scientists constantly analyze new samples and review data from customers who opt-in to share metadata to uncover new and emerging attack behaviors and trends. This data continually feeds supervised machine learning algorithms, which are shared with all Vectra customers worldwide.

Unsupervised machine learning

Unsupervised machine learning shifts the focus from global sets of data to learning what is normal among the unique characteristics of the environment that is being protected. These models focus on behavioral anomalies and accurately detect a wide range of attack techniques.

For example, pass-the-hash techniques have been an essential network-based attack tool for years. Each year, network and PC vendors roll out new protections that are designed to stop the current crop of pass-the-hash tools. And every year, attackers release new tools to defeat those controls.

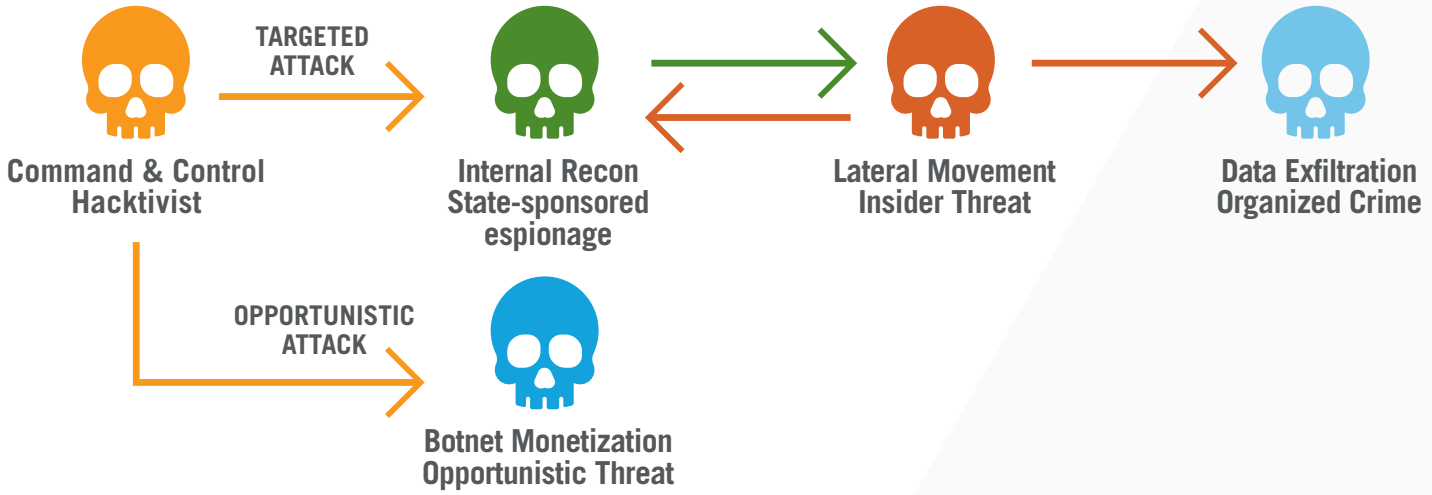

Cognito constantly monitors user and administrative behavior and tracks the user and admin accounts and services that are requested by different devices, including the underlying protocols used to manage servers and cloud environments. When an attacker attempts to pass-the-hash, Cognito identifies the behavior equally across the newest and oldest forms of pass-the-hash.

Detecting threats in encrypted traffic without prying

SSL/TLS and other types of encryption pose a challenge for most security products. But by focusing on malicious actions instead of malicious payloads, Cognito identifies active threats in encrypted traffic without decrypting the traffic.

Cognito continually reveals the underlying purpose of traffic, even when the payload is not visible. This is a critical development because it allows security teams to protect without prying.

Cognito identifies active threats without decrypting the traffic by focusing on malicious actions instead of malicious payloads.

Cognito even finds attackers who tunnel hidden communications within an SSL-encrypted web session. By analyzing tiny fluctuations in protocols like HTTPS, HTTP and DNS, Cognito reveals when additional layers of communication are hidden within.

This is a vital set of capabilities. Vectra's own research has found HTTPS to be the most popular protocol for these hidden tunnels. And by detecting threats in encrypted traffic without decrypting it, Cognito mitigates attacks with no performance penalty or privacy concerns associated with decryption.

Applying intelligence to all phases of attack

Sophisticated attacks are long-term strategic operations that naturally progress through multiple phases. Unlike old models that emphasize detecting the initial compromise, NDR must detect all phases of an attack.

It is important to detect in-progress threat behaviors in every phase of the attack lifecycle

Command-and-control and remote access

Attackers depend on command-and-control and remote access tools to orchestrate and advance their ongoing threat activities. These attacks are only possible if the remote attacker maintains ongoing control of devices inside the network.

Many security solutions rely on signatures and reputation lists to identify command-and-control traffic, but they have severe limitations. Command-and-control signatures work well for large, well-known botnets. But they are easily evaded by attackers who customize their command-and-control infrastructure and use each variant for only one target organization.

Command-and-control and remote access tools are used to orchestrate and advance attack activities.

Cognito identifies a wide range of command-and-control behaviors, including attempts to imitate browser behavior, use of hidden tunnels, peer-to-peer communication, malware updating as well as a broad variety of anonymization techniques such as TOR.

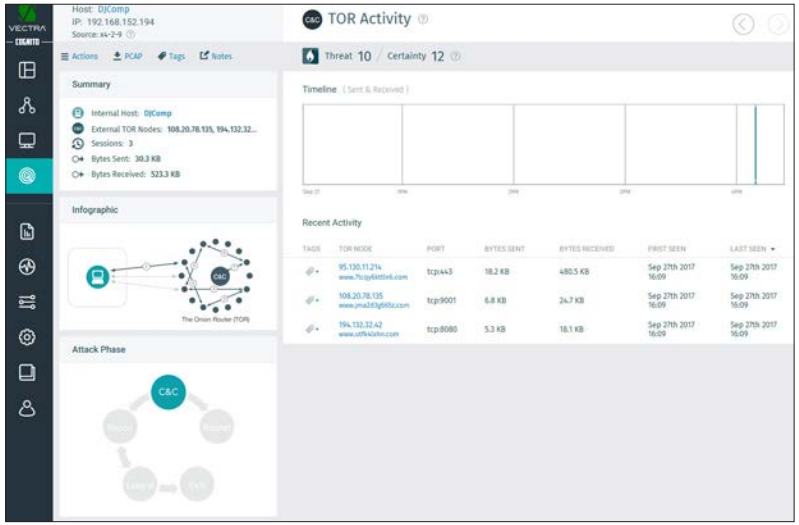

Command-and-control detection of TOR anonymization activity

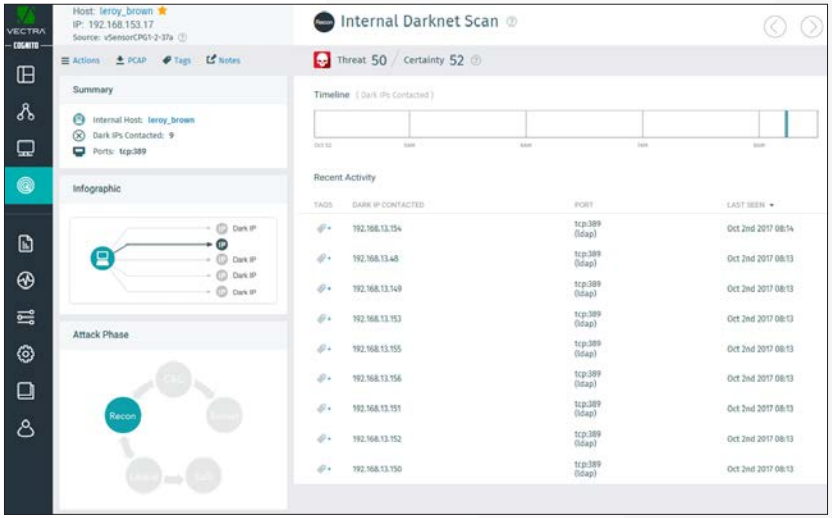

Reconnaissance detection of an internal darknet scan

Likewise, Cognito identifies all types of external remote access tools that attackers use to directly control infected hosts. As with all attack behaviors, Cognito reveals this behavior generically and even if traffic is encrypted. This ensures that even the newest variants of malware are always detected.

Internal reconnaissance

After attackers gain access to a network, the attack processes begin anew. The initial victim machine usually doesn't contain the most valuable data in the network. As a result, attackers will learn the local network environment and identify other hosts and segments to exploit.

Cognito identifies reconnaissance behaviors, even if attackers take a low-and-slow approach to map out the network. In addition to identifying reconnaissance being performed inside the network, Cognito scans individual host machines that are targeted for attack.

Lateral movement

The most crucial phase of a cyberattack involves lateral movement. The ability to spread laterally inside the network provides attackers with places to maintain persistence and enables them to dive deeper as they progress toward key assets.

Lateral movement takes one of two forms. Attackers will spread malware inside the network from host to host or steal credentials from victims to access critical network resources.

Cognito exposes both forms of lateral movement. By monitoring all internal traffic, the patterns of a host that is spreading a malicious payload to other hosts are quickly recognized. Again, Cognito reveals this spreading behavior without inspecting or analyzing the payload.

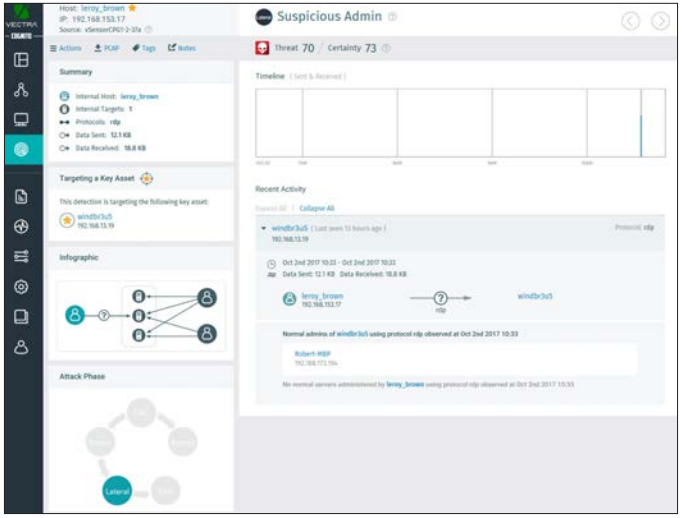

Lateral movement detection of suspicious admin activity

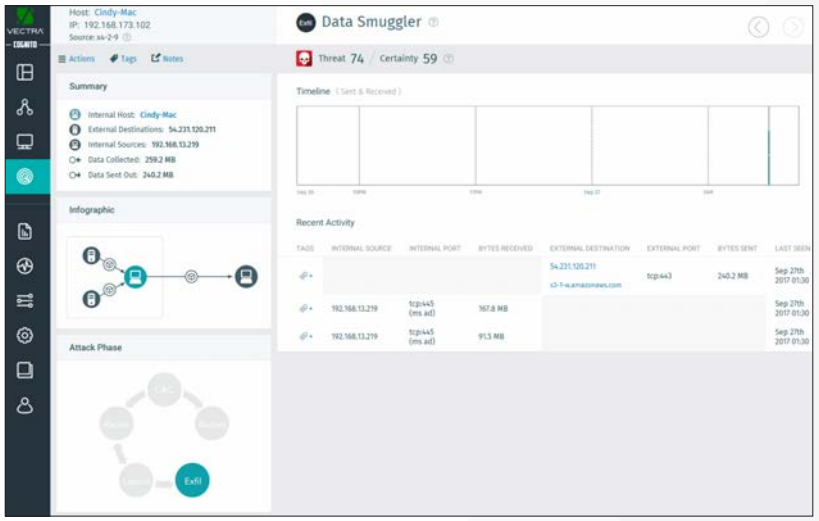

Exfiltration detection showing a data smuggler

Lateral movement involves spreading malware from host to host or stealing user credentials to access vital network resources.

In the case of stolen credentials, Cognito constantly monitors the internal Kerberos infrastructure to identify signs of theft or credential re-use. This capability reveals very subtle attacks, even when no malware is involved.

Data acquisition and exfiltration

The final phase of attack involves acquiring data and sending it back to the remote attacker. Cognito monitors the network for devices that are acquiring and sending data at an abnormal rate.

Additionally, the exfiltration process requires attackers to stage data for aggregation. The data is typically moved to areas of the network where uploading draws less suspicion. Automatically and in real time, Cognito connects the dots and recognizes when data is being staged and prepared for transfer.

Detecting attacks, not events

IDS is notorious for generating mountains of event logs that require an extensive investment in resources and time to investigate. Security teams are overwhelmed by a steady stream of alerts that turn out to be false positives and simply ignore a majority of them.

Instead of relying on individual events or detections, Cognito uncovers in-progress cyberattacks inside a network. Detections are correlated to the hosts under attack, and each is scored and prioritized according to the highest risk. Hosts with detections are plotted in the Cognito Threat Certainty Index™, which instantly reveals hosts at the center of an attack.

Detections are correlated to the hosts under attack, and each is scored and prioritized according to the highest risk.

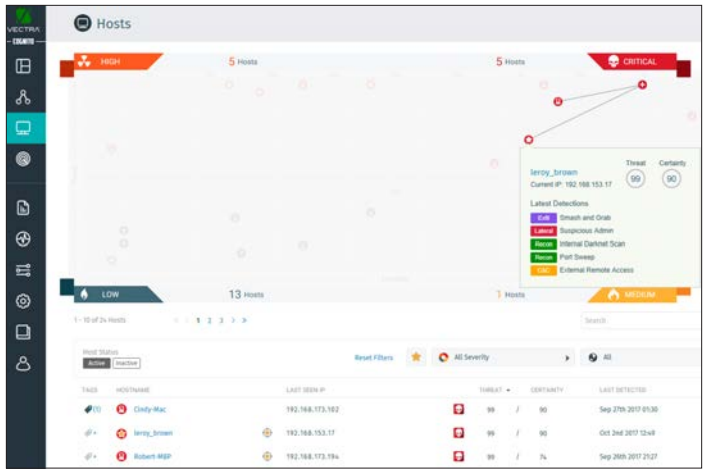

The Cognito Threat Certainty Index

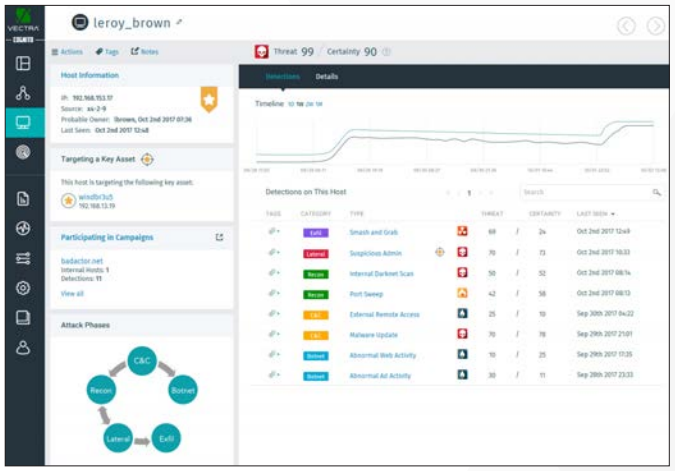

A view of host details shows the progression of attacks over time

Tracking the attack progression

Scoring goes beyond aggregating the number of detections tied to a host. A host's threat score increases as Cognito observes multiple phases of an attack. For example, a host associated with reconnaissance, lateral movement and command-and-control detections is prioritized above a host that simply shows a large volume of botnet monetization behaviors.

Security teams can also track the progression of an attack over time. A view of host details shows a history of all detections as well as an hourly analysis of threat and certainty. If security teams need to dig deeper, they can instantly access metadata from packet captures for any Cognito detection.

Not all hosts are created equal

Although all hosts on a network are important, they are not all equal. Cognito takes this into account and gives the option of marking key assets. This allows security teams to easily track and prioritize events in critical areas, such as servers that contain data in the scope of PCI DSS.

Continuously monitor privileged access

A traditional, access-based approach to zero trust relies on one-time gating decisions that use a predefined list of privileged identities. This approach is fundamentally flawed when cyberattackers have already obtained credentials or have escalated privileges.

In the cloud and other environments that lack traditional security boundaries, the need for continuous real-time assessment of user, host and service privilege levels is especially pronounced.

In addition, the complexity of access management makes it prone to misconfigurations. Continuous monitoring and alerting for unusual privileged access is a critical requirement for modern NDR solutions.

Conclusion

IDS continues to lose its edge in detecting intrusions as modern cyberattackers gain momentum using more evasive and sophisticated methods to spread rapidly throughout the network. This leaves security teams without the means or visibility to identify threats that pose tremendous risk to their organizations.

Cognito, with its AI-powered cyberattack detection capabilities, is the ideal replacement for today's IDS products that cannot block contemporary cyberattacks and cannot detect hidden attacker behaviors inside your network.

It's time to jettison the moth-eaten limitations of IDS and concentrate on detecting and mitigating active threats inside the network – from users to IoT devices to data centers and the cloud – before attackers have a chance to spy, spread and steal.

For more information please contact a service representative at sales-inquiries@vectra.ai.

Email info@vectra.ai vectra.ai

Vectra® protects business by detecting and stopping cyberattacks.

As a leader in network detection and response (NDR), Vectra® AI protects your data, systems and infrastructure. Vectra AI enables your SOC team to quickly discover and respond to would-be attackers —before they act.

Vectra AI rapidly identifies suspicious behavior and activity on your extended network, whether on-premises or in the cloud. Vectra will find it, flag it, and alert security personnel so they can respond immediately.

Vectra AI is *Security that thinks*®. It uses artificial intelligence to improve detection and response over time, eliminating false positives so you can focus on real threats.