White paper

VECTRA®

How Vectra replaces IDS and enables organizations to detect intrusions again

VECTRA

TABLE OF CONTENTS

Introduction

Intrusion detection systems (IDS) have been a mainstay of information security for decades. Over the years, IDS technologies were gradually subsumed by intrusion prevention systems and today they are known collectively as (IDS/IPS). In most security circles, IDS is simply thought of as IPS that is deployed in a passive or listen-only mode.

However, this mode of thinking is glaringly outdated. As persistent cyber attacks and network breaches become more common, the need for intrusion detection is higher than ever before. However, IDS is not living up to its name because intrusions are happening and they are not being detected.

To address this problem, it's important to remember an inherent order of operations. Detection must come before prevention. You can't block what you can't detect. As such, it's not enough for IDS to simply be a defanged version of IPS. IDS must again be on the innovative leading edge of security that detects threats that other approaches can't.

This paper lays out a new model for network intrusion detection based on the Vectra® Networks platform. It includes a brief look back at the history of intrusion detection and how it led to some of the limitations we face today.

The remainder of the paper focuses on how to move intrusion detection forward. It explains the unique challenges of detecting modern threats and proposes 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 ways 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.

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.

By the 1990s, the first commercial network-based IDS arrived. Armed with the ability to sniff packets, they began a shift from behavioral to misuse models that looked for specific signs of an attack within packets or network sessions. This allowed a faster, more automated approach to finding threats as long as they had known indicators.

The shift from IDS to IPS

Network-based IDS also marked an architectural shift. Instead of an internal focus, which was the norm with host-based approaches, network IDS shifted to the perimeter. As the Internet grew, more and more threats were coming from outside the walls of the enterprise, and security teams needed a way to defend their applications and assets.

> *As the Internet grew, more and more threats were coming from outside the walls of the enterprise.*

While network firewalls largely played the role of traffic cop by separating the internal network from the outside world, they were not up to the task of finding and stopping attacks. As IDS became faster and more reliable, it was becoming possible to shift from detection to prevention. The IPS was born, and steadily became a critical and standard layer of the network perimeter.

IPS and the hidden impact on intrusion detection

Since the early 2000s, the evolution of IDS/IPS has been overwhelmingly driven by the demands of IPS. There are good reasons for focusing on IPS: The Internet harbors a virtually unlimited number of threats and security teams needed a way to block the vast majority them automatically. Over the years, IDS/IPS evolved into stand-alone, in-line appliances and became part of next-generation firewalls, UTM or blade architectures.

Speed is king

The in-line approach to detection and prevention comes at a price. For in-line security devices, networking concerns must come before security concerns. High speed, high throughput, and low latency are non-negotiable requirements. These demands have led to constant trade-offs that have sacrificed detection capabilities for speed.

First, detections 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.

This requirement has locked the vast majority of IDS/IPS in a signature-reliant approach to detection. 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.

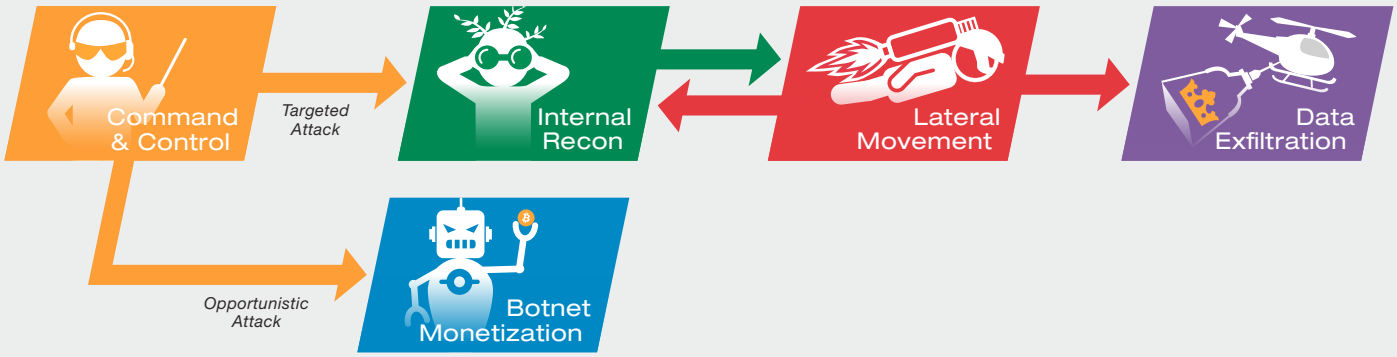

Vectra detects active cyber threats in every phase of the attack kill chain

Likewise, the detection logic of IPS must consume an absolute minimum of memory. Most IPS products can't enable all their signatures for performance reasons, which forces security teams to pick and choose certain signatures to find the right balance of security and speed.

The more memory a signature requires, the fewer signatures can be run. This means IPS signatures must be fast and the detection logic must have a very short time horizon. Detections based on long-term analysis of a series of sessions or user behavior over time is not possible.

Putting detection first

IPS and IDS are no longer deployment options of the same technology. They 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.

As threats grow more sophisticated and intrusions become more common, a new generation of IDS is required that once again makes detection the top priority. The IDS must be defined not as a subset of IPS features but as the intelligence that detects the wide variety of network threats that IPS cannot.

> *A new generation of IDS is required that once again makes detection the top priority.*

This approach logically separates the unique demands of intelligence and enforcement without compromising one or the other. The next-generation of IDS becomes the metaphorical brain while IPS remains the muscle behind enforcement.

To meet today's challenges, IDS must advance the state of the art for threat detection, which requires a blend of new detection strategies and techniques. It should include the ability to detect threats even if no malware or exploit is used.

IDS must evolve to detect modern threats wherever they hide. This requires extending visibility past the network perimeter to internal segments where intruders lurk. IDS must also have an extended time horizon to identify the progression of an attack.

Modern threats and the need for true intrusion detection

Attackers have adapted to the perimeter security model in two fundamental ways – perimeter evasion and perimeter avoidance. Evasion involves getting threats past the perimeter without detection, while avoidance finds avenues into the network without crossing boundaries. We will look at some of the more common approaches to each and how they can be addressed.

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 it. 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 let attackers 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 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.*

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 data at the perimeter.

Perimeter avoidance

In addition to techniques that smuggle cyber threats across the perimeter, attackers have learned to avoid the perimeter altogether. By infecting users while at home or outside the perimeter, threats can be carried in the front door on a victim's device.

Mobile devices provide logical and physical paths around the perimeter. Mobile devices with LTE or 4G data connectivity have an easy path to the Internet and can introduce serious risks that never have to cross the network perimeter. It's an invisible conduit that attackers love to use to get inside networks.

Shifting to the inside of the network

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.

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 provides 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 tools for attackers that are allowed by network policy.

Defining the next-generation of IDS

To detect active network intrusions, IDS must revisit concepts from its past as well as introduce a variety of completely new techniques. Attackers that penetrate the network perimeter are fairly sophisticated. No single technique or approach will ensnare an intelligent and creative adversary, so the modern IDS must be flexible and adaptable.

The importance of internal visibility

As the name implies, intrusions occur largely on the inside of the network. Clearly, IDS has little hope of detecting threats if it can't see where the majority of action is taking place.

> *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 important for IDS to monitor internal network segments as well as the Internet boundary. Internal reconnaissance, lateral movement, unauthorized data access and staging only occur inside the network. Without internal traffic visibility, it's nearly impossible to detect these attack phases.

For IDS, this means moving deeper into the network. Instead of deploying IDS at ingress and egress points, it should be deployed where it provides the greatest traffic visibility. While every network architecture is different, passively monitoring traffic at the core switching infrastructure in an office is a good first step.

It's also important to account for network segments that house key assets and critical data. An IDS must be modular and flexible enough to support these areas efficiently and provide full visibility.

The network provides a trusted perspective

Historically, endpoint security and host-based IDS were employed to detect threats inside the network. While these technologies are important parts of any security strategy, they do have significant limitations.

First, it's unrealistic to ensure that all endpoints are protected. Users connect a wide variety of devices on a daily basis that often evade security. Even if a security team knows about every device, the variety of operating systems, device types and updates make it nearly impossible for host security applications to support all devices.

The challenges of Internet-of-Things (IoT) devices compounds this problem even further as printers, phones, alarm systems, thermostats and hundreds of other IP-enabled devices enter the modern network.

While installing security software on every host can be logistical nightmare, it also suffers from a logical problem. The age-old challenge of endpoint protection requires a security application to defend a potentially compromised host.

A new or powerful rootkit can gain control over a device and in the process subvert the security application itself. Other malware can target the BIOS or the Baseboard Management Controller, which can give attackers control below the level of the operating system.

These examples point out the circular problem of host-based security. It's impossible for host-based security to be independent of its host. If and when the host is compromised, the security program itself is a guest on a contaminated vessel. This falls back into the same problem of prevention-only security – after prevention fails, attackers have the upper hand on the device.

> *Passively monitoring network traffic offers an impartial way to observe the real behavior of hosts.*

Passively monitoring network traffic offers an impartial way to observe the real behavior of hosts. A host's nefarious actions will ultimately show up on the network, particularly those that involve an attacker spreading, accessing data or communicating with the outside world.

Just as important, due to its passive deployment, IDS can sit above the fray. This third-person objectivity breaks the circular problem of trust found in host-based and even some in-line security solutions.

New models of threat detection

Moving from payloads to behaviors

To detect more sophisticated attacks, IDS must move beyond the realm of traditional signatures. One of the underlying limitation of signatures is that they typically search for malicious payloads.

This can be byte-level signatures that identify known exploit code within a packet or a hash-based signature that matches known malware. Similarly, signatures can attempt to match IP addresses and URLs to blacklists or reputation lists of bad sites and addresses that have been seen in the past.

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. This is due in part because the detection model is looking for specific payloads.

To move beyond his challenge, detection models should focus on identifying the underlying malicious behavior. 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 because while attackers can easily put on a new coat of paint to avoid signatures, they still need to accomplish similar goals. Attackers have a near-infinite supply of tools to help them to spy, spread, and steal within the network, but they must perform the same fundamental actions and behaviors to succeed.

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.

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 malware that updates the binary should be detected, whether it occurs via webmail, Twitter or Dropbox.

The science of detection

The Vectra® Networks approach to detection is based on the direct analysis of traffic to reveal the fundamental behaviors at the heart of cyber attacks.

By combining data science, machine learning and behavioral analysis, Vectra identifies the underlying purpose of traffic and reveals malicious behaviors, independent of applications and even when the traffic is encrypted. This approach reveals the key actions that an attacker must perform in order to succeed.

Unlike data analytics products that analyze logs from other devices, Vectra applies algorithmic models directly to network traffic to reveal underlying attack behaviors that are not visible based on logs or flow analysis.

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, Vectra identifies the unique pattern of an outsider who is controlling a machine inside the network. We will explain these concepts in more detail in the next sections, which delve into the differences between Supervised Machine Learning and Unsupervised Machine Learning.

> *Vectra applies algorithmic models directly to network traffic to reveal underlying attack behaviors.*

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 crucial 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 Vectra algorithms that detect underlying malicious behavior in network traffic.

While global intelligence is certainly required, 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, Vectra models learn and detect based on times ranging from seconds to weeks.

Supervised Machine Learning

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

By identifying common underlying behaviors, Vectra breaks the cycle that plagues signature-based solutions. Vectra models know the unique patterns of command-and-control servers that guide internally-infected hosts. These behaviors are distinguishable from normal traffic and recognized across malware families, even when commands are carried inside valid trusted application protocols.

> *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. The Vectra model 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.

Vectra constantly monitors user behavior and tracks the user accounts and services that are requested by different devices. When an attacker attempts to pass-the-hash, Vectra 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 particular challenge for most security products. But by focusing on malicious actions instead of malicious payloads, Vectra identifies active threats without decrypting the traffic.

Vectra algorithms continually reveal 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.

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

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

This has become 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 without decrypting traffic, Vectra 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, modern IDS must detect all phases of an attack.

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.

Vectra 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.

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

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

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.

Vectra detects 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, Vectra scans individual host machines that are targeted for attack.

Lateral movement

The most crucial phase of a sophisticated cyber attack 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.

Vectra detects both forms of lateral movement. By monitoring all internal traffic, Vectra recognizes the patterns of a host that is spreading a malicious payload to other hosts. Again, Vectra detects this spreading behavior without inspecting or analyzing the payload.

In the case of stolen credentials, Vectra 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.

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

Data acquisition and exfiltration

The final phase of attack involves acquiring data and sending it back to the remote attacker. Vectra 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, Vectra 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 the vast majority of them.

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

Instead of relying on individual events or detections, Vectra detects in-progress cyber attacks 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 Vectra Threat Certainty Index™ which instantly reveals hosts at the center of an attack.

Tracking the attack progression

Scoring goes beyond aggregating the number of detections tied to a host. A host's threat score increases as Vectra 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 Vectra detection.

Not all hosts are created equal

Although all hosts on a network are important, they are not all equal. Vectra 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.

Conclusion

IDS continues to lose its edge in detecting intrusions as modern cyber attackers gain momentum using more evasive and sophisticated ways 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.

Vectra, with its real-time automated threat hunting capabilities, is the ideal replacement for today's IDS products that cannot block contemporary cyber attacks and cannot detect 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 using Vectra – before attackers have a chance to spy, spread and steal.

VECTRA®
Security that thinks®

Email info@vectranetworks.com **Phone** +1 408-326-2020
www.vectranetworks.com