CAPSULE8

Why IDS is Ineffective for Linux Production Environments

A Capsule8 Technical Primer

Nine Reasons and Zero Days

In the alphabet soup that is a traditional cybersecurity architecture, intrusion detection systems (IDS) hold a prominent spot on the mantle. IDS are broadly recognized as an important component of a cybersecurity strategy; one of many tools that make it more difficult for an adversary to inflict harm on an organization.

The threats companies face become even more pressing in the context of production environments. Production systems – those housing customer data, IP and other critical information – must be protected holistically. As the heart of a business, attacks that impact production infrastructure have the potential to cripple organizations, including the potential of fines, such as those Google has seen of late related to GDPR.

Few would argue the necessity of being able to detect intrusions into the production infrastructure. IDS and other "good hygiene" technologies – such as firewalls, antivirus and strong authentication – play a central role in helping companies to fend off threats. An IDS, in particular, helps monitor a production environment for unusual or malicious activity, either at the host-level (host IPS – HIDS) or network-level (network IPS – NIDS).

In this context, an IDS may be considered a "911 operator" of cybersecurity. These technologies identify threats and summon resources to take action, but are not chartered with doing so themselves. Whether an IDS alert is delivered to an IT administrator or a central management platform (such as a SIEM), the technology is designed to provide visibility after a suspicious or malicious activity has already occurred.

In short, IDS plays a core role in ensuring fundamental cyber monitoring capabilities are in place. While IDS plays a key role in a defense strategy, the technology – most often delivered via a hardware appliance – faces additional challenges in the context of production infrastructure. Much more is needed to ensure companies are effectively protecting themselves and their customers, particularly in the context of increasingly complex production environments that are the norm today.

Managing the external environment becomes even more challenging in the context of the ongoing digital transformations underway at many organizations. In a drive to gain scale and efficiencies, companies are leaning more heavily on cloud-based capabilities. Organizations are evolving and modernizing their production environments with technologies like cloud, microservices and containers, and are more often mixed with both cloud and on-premises infrastructure and applications. This creates a changing attack surface that conventional security solutions such as IDS simply cannot address. And with vulnerabilities such as Meltdown and Spectre, legacy Linux infrastructures are also up against inadequate protection caused by low visibility and poor detection. As crucial as IDS has been in the past, they are eventually going to die out as the technology landscape continues to evolve. They'll continue to get less effective, pushing detection to the machines that need protection as the need for real-time and zero day attack detection grows. This new reality is forcing cybersecurity leaders to take a long and hard look at the strategies and capabilities they deploy to secure personal data and critical business information.

The reasons IDS is not effective can be boiled down to nine main points:

1. IDS CREATES BOTTLENECKS IN A CLOUD-BASED INFRASTRUCTURE

It's not a secret that a lot of infrastructure is moving to the cloud. Even Capital One has moved 60% of theirs already. FINRA, the Financial Industry Regulatory Authority, has moved 75% of their infrastructure to Amazon Web Services. Organizations aren't willing to pay traffic cost or latency to hairpin out to an IDS, and vendors will try to provide "virtual appliances" within the cloud, which is an unnecessary bottleneck.

2. MANAGING HYBRID ENVIRONMENTS IS TOO DIFFICULT FOR IDS

Hybrid deployments (i.e., partially on premise and partially in the cloud) will undoubtedly be a fact of life in the enterprise for a long time to come. That means, security teams must provide solutions for both kinds of environments. Today, while we're still early in the adoption curve, security organizations are willing to implement different solutions to protect their cloud infrastructure and their internal infrastructure. Yet, it will become a burden—more costly, and more cumbersome to manage, so people won't do it forever.

3. IDS CAN'T SEE CONTAINERS

The DevOps movement is pushing toward microservices and containerization quickly. When multiple containers live on the same machine they operate concurrently and communicate frequently. That communication doesn't go over the network and can never be seen by an IDS or appliance — not even a virtual appliance. As technology teams deploy multiple containers on a single server, communication between these is confined to the machine. Without network-based communication, a NIDS is unable to view communication and, in turn, any potential threats that may be resident. In addition, it's important to recognize that appliances rely on IP or host names for monitoring, but in a containerized world, containers tend to be short-lived and many containers can share an IP address. That lack of visibility means appliances are far less effective at detection for modern production environments.

4. IDS APPLIANCES DON'T ACTUALLY DETECT

In most enterprises, IDS appliances are usually sitting off to the side, looking at a copy of network traffic, not the actual network traffic. Organizations do this for many reasons: so that appliances eliminate unnecessary latency to network traffic, and so that they don't become a single point of failure, for instance if they get flooded or have a bug. Worst of all, since it's so hard to get high quality signal from network traffic at scale, appliances generate many false positives, which are a huge disaster for automatic response. And while traffic in the cloud can still be split off, it isn't easy and it can come at a price.

5. IDS ARE EASY TO CIRCUMVENT

While old-school devices (e.g., traditional Intrusion Prevention Devices) obviously sacrificed accuracy for speed, today's more sophisticated appliances do a lot of processing on the data they see, so that they can give vastly better results, with far fewer false positives. Because they rely on emulation, an attacker has many options to detect and circumvent the emulation. Eventually, detection will move to the systems being protected, where there is no need to emulate, and there's a much greater ability for security software to thwart an attack.

6. IDS CANNOT AUTO-SCALE

Generally, the IDS with the best detection consumes the most resources. Even with a high-speed appliance, it's generally not difficult to overload them. Once an attacker manages that, they can sneak malicious traffic through undetected. As companies embrace infrastructure that can auto-scale their applications, they will want to auto-scale their protection to improve their zero day attack detection, instead of failing open.

7. IDS ARE TOO MUCH WORK FOR TOO LITTLE VALUE

Perhaps the single biggest problem that IT Security organizations wrestle with is that they're drowning in alerts. Generally, that's the case even AFTER all the raw data coming from around the network goes through a best-of-breed correlation and analysis engine. This problem is due to the horrible signal-to-noise ratio in security appliances, which is going to get even worse. Instead of hiring more analysts or letting more and more alerts slip through the cracks, companies will look for detection approaches that provide much lower noise, which again pushes detection away from an IDS or appliance solution.

8. IDS CAN'T VIEW THE DATA REQUIRED FOR REAL-TIME ATTACK DETECTION

In the United States, Edward Snowden's disclosures were a wake-up call to those who thought encrypting data across the Internet's backbone wasn't particularly important. Many people used to think it was unlikely that anyone would have the means and desire to listen. It turns out, they were wrong. (Note that, in many other countries, the Government is quite explicit about this kind of access to Internet traffic). The standards community, in protecting against nation state attacks, are making it impossible for security appliances to rely on what is essentially a decryption back door, meaning companies will have to incur huge expense to provide both data privacy and give security appliances visibility to do detection. This is already on the horizon, with the forthcoming version of TLS, which secures every HTTPS connection.

9. IDS CANNOT RESPOND TO ATTACKS

Remember, the role of an IDS is not to prevent an attack from occurring, but to identify and alert administrators and other systems (e.g., a SIEM) that a threat has been identified. In this context, it is also critical that an organization has the human power in place to act upon information delivered by an IDS; without follow-up action, the flags raised by an IDS are for naught.

And Then There Were Zero

ZERO DAY ATTACKS SKIRT INTRUSION DETECTION SYSTEMS

Historically, IDS have relied on a library of signatures, which hold a repository of known threats. These signatures enable the IDS to identify a previously-seen attack vector, allowing the system to quickly alert an administrator. While responding to known threats is necessary, reliance on signatures carries a significant disadvantage: the inability to fend off zero-day attacks.

Consider a test executed against Snort, an open source NIDS designed to identify suspicious network activity, as reported in an article published by the Institute of Electrical and Electronics Engineers. The article reports that “while the actual detection rate of the tested zero-days’ was 17%, this number does not consider the possibility of false alarms or signature evasion techniques … an overall of 48.8% of all alerts can be considered effective. Thus, a conservative estimate on the overall detection rate by Snort for zero day attacks is 8.2%.” In short, the article reports that while Snort detection of zero-days is greater than zero, “SNIDS is not able to provide complete detection of either known attacks or zero-days’.”

It’s clear, there is an inherent disconnect between the traditional characteristics of an IDS and the requirements for identifying and responding to zero-day threats. The question becomes, though, how organizations manage the challenge of threats that have never been seen – and, therefore, lack signatures that support the role of the IDS.

As organizations look to address the challenge of both detecting and responding to zero-day attacks across production environments,encompassing containers, virtualization and bare metal, it's clear that IDS alone will not answer the call.

Enter Capsule8: Attack Detection for Modern Linux Infrastructure

As companies look toward solutions for both detecting and responding to attacks, particularly in the context of critical production infrastructure, Capsule8 offers unmatched cybersecurity capabilities. Capsule8 is the industry's only real-time, zero-day exploit detection platform ***purpose-built for Linux production environments*** – whether containerized, virtualized or bare metal. Capsule8 massively reduces security operations' workload by automatically detecting and shutting down exploits as they're happening – without adding any risk to production infrastructure.

UNLIKE AN IDS, CAPSULE8:

Detect exploits in real time – Capsule8 uses distributed, streaming analytics combined with high-fidelity data that detects attacks in the instance they're attempted. This real-time approach allows our customers to respond to attacks before they have costly consequences - most often before the attack takes hold. While most security products focus on Indicators of Compromise (IOC) - which indicate reactive/retrospective awareness - Capsule8's approach is proactive, detecting Indicators of Attack (IOA). This active awareness positions every organization to better control the scope of risk and impact and allows for less configuration churn.

Uses a multi-layer approach to detection ensuring high-fidelity alerting – Capsule8 has unparalleled, system-level detection that is continuously updated to uncover the latest zero-day attacks. Our strategies include highly technical methods for detecting indicators of common exploitation techniques, while still providing flexible policy-based detection (such as file integrity monitoring). These multi-layered detection strategies are devised by a team of security researchers and data scientists unrivaled in the industry.

An attacker will, in the course of attacking a system, perform a series of actions that will trigger individual strategies. As a result, a customer's typical flood of alarms is reduced to a trickle of alerts around actual exploits. Capsule8 also developed "kernel landmines," triggers placed in a running Linux kernel associated with a process that shouldn't normally be touched by regular authorized processes and application usage. This allows Capsule8 to monitor places in the kernel that are possible windows into exploit behavior.

Adds no risk to production – Capsule8 uses a variety of techniques to ensure that the solution will not have an undue impact on production, and will be easy for your Ops teams to manage. To ensure minimal performance impact to hosts and networks, Capsule8 employs a resource limiter that enforces hard limits to system CPU, disk and memory, with an intelligent load-shedding strategy. This is fully adjustable so users can choose their ratio of detection accuracy to CPU usage. Everything at Capsule8 is engineered, from the start, to have a minimal impact on production.

Provides clear and actionable attack intelligence – Capsule8 offers a level of transparency that makes it easy to determine why alerts fire and what an attacker does after an attack lands. In addition to observability hooks, we have context in the form of metadata awareness and allow arbitrary metadata creation and policies to be made along with it.

Is tunable to your environment – Capsule8 makes it easy to implement complex and customizable policies that vary not just by machine, but by far more granular system-level items, helping to minimize false positives.

Automates attack response – Capsule8 helps customers disrupt attacks in real-time, as they're happening. For instance, customers can strategically (and automatically) kill attacker connections, restart workloads, or alert an investigator, immediately upon initial detection. Our automated disruption minimizes the attack impact, preventing damage as well as the need for any unwanted manual investigation and cleanup, enabling stronger and more resilient systems. This is possible because of Capsule8's correlation engine -- a critical piece of detection engineering. Every possible automation and correlation is fired before a human is brought into the loop, and if human is engaged, they receive the proper context and recommended action. Capsule8 developed creative signal processing and stochastic methods, precisely aimed to produce high quality and high-signal insight that is meaningful to analysts and developers. This makes real-time, automated analysis and response possible.

Is built for production, providing one solution for cloud-native & legacy environments – We work where you work. Capsule8 provides seamless, easy-to-deploy detection across a wide variety of Linux versions, be it public cloud or data center, containers, virtual machines or bare metal. We protect all major Linux orchestrators, including Kubernetes, Docker, and CoreOS and configuration management tools such as Puppet and Ansible.

CAPSULE8

About Capsule8

Capsule8 is the industry’s only high-performance attack protection platform that’s safe for the busiest workloads, on the busiest networks. Capsule8 delivers continuous security across your entire Linux production environment — containerized, virtualized and bare metal — to detect and disrupt attacks as they happen.

Learn more at www.capsule8.com.