GIAC
CERTIFICATIONS

Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permited without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Network Monitoring and Threat Detection In-Depth (Security 503)"
at http://www.giac.org/registration/gcia

Intrusion Prevention System Signature Management Theory

GIAC (GCIA) Gold Certification

Author: Josh Levine, josh@cantreach.me
Advisor: Rajat Ravinder Varuni
Accepted: 2/4/2019

Abstract

The intrusion prevention system (IPS) serves as one of the critical components for a defense-in-depth solution. IPS appliances allow for active, inline protection for known and unknown threats passing across a network segment at all layers of the OSI model. The employment, tuning, and upkeep of signatures on an IPS may lead to a negative impact on production traffic if not properly maintained. This document serves as baseline guidance to help shape the development of an organizational IPS signature management policy. Concepts are presented to address the lifecycle of an IPS signature from employment to expiration. Through proper maintenance, placement, and tuning of signatures, an unwanted impact to network traffic can be kept to a minimum while also achieving an optimal balance of security and network performance. By understanding the tenants of effective IPS signature evaluation, employment, tuning, and expiration, organizations can maintain an acceptable network security posture along with adequate levels of network performance.

© 2019 The SANS Institute, Author Retains Full Rights

© 2019 The SANS Institute Author retains full rights.

1. Introduction

Both the intrusion prevention system (IPS) and intrusion detection system (IDS) serve as fundamental components of a defense-in-depth solution. Relying on a combination of heuristics, anomaly-based, and behavioral-based detection mechanisms, and signature sets designed to detect potentially harmful traffic, these systems provide active protection and detection for various threat vectors occurring on a given network segment. The management and tuning of signatures can have a direct correlation to network performance, reliability, and overall threat posture.

For this paper, the focus will be on generic signature management terminology as implemented on the active intrusion prevention system (IPS) devices. Any concepts covered are meant to apply in a vendor-agnostic manner. Actionable recommendations are provided based upon established best practices in the cybersecurity industry. These recommendations are designed to shape the development of organizational guidance supporting operational signature management and refinement.

It is important to note that this paper is not intended to describe hard and fast rules which must be followed for successful signature management or IPS/IDS operation. This paper is also not intended to provide rules and guidelines that will work for all organizations. Instead, the purpose is to assist in guiding decisions made with regards to IPS/IDS signature management and tuning procedures. Organizational requirements will vary, and this guidance should be adjusted based on specific operational requirements.

2. Key Terminology

2.1. Overview

A common lexicon has been provided to establish a baseline vocabulary for terminology used throughout this paper. The below terms are offered to level-set and provide a standard terminology for all readers.

Josh Levine, josh@cantreach.me

© 2019 The SANS Institute Author retains full rights.

© 2019 The SANS Institute, Author Retains Full Rights

2.2. Terminology

Action: A act taken against specific traffic flows, e.g., allow, pass, block, transform.

Anomaly: A behavior or action occurring outside of a defined behavioral or characteristic-derived baseline (Chandola, Banerjee, & Kumar, 2009).

Authorizing Official (AO): The individual within an organization that accepts responsibility for the overall security and risk posture of a given network enclave.

Common Vulnerabilities and Exposures (CVE): A list of entries about a publicly disclosed threat or vulnerability in a product. Each CVE is linked to an identification number used to reference a given vulnerability, or a list of vulnerabilities, in the associated product (Mitre, 2018).

Deep Packet Inspection (DPI): The process of analyzing traffic through all layers of the OSI model (layers one through seven). Deep packet inspection is the core capability that separates an IPS from a traditional firewall. DPI is also a capability of next generation and application-aware firewalls.

Defense in Depth: The use of a multi-layered approach to security that encompasses all aspects of an organization's security posture. This methodology relies on the protection and mitigation mechanisms spanning the security spectrum to include: the human layer, transport layer, network layer, application layer, and host layer. The overlapping of controls across these layers will ensure adequate protection should one measure not actively mitigate an identified risk. (Barnum, Gegick, & Michael, 2005)

© 2019 The SANS Institute Author retains full rights.

Detect:	The ability to identify a given characteristic or behavior.
East/West Traffic:	Traffic internal to a given network enclave (host <-> server, server <-> server, host <-> host). (Scarfò, 2011)
Failback Mode:	This mode turns the IPS into a passive transport device, bypassing the signature and traffic processing engine in its entirety to maintain an operational capability for network traffic.
Flow:	A set of characteristics for network traffic based upon a standard five-tuple for a given network conversation. The five-tuple consists of source and destination IP addresses, source and destination ports, and traffic directionality (A->B, B->A, A<->B).
Information Assurance Vulnerability Management (IAVM):	The process by which an organization manages their hardware and software configuration baselines to minimize risk and control vulnerabilities. The IAVM process includes steps designed to rapidly detect and respond to vulnerabilities identified through routine scanning or an industry alert such as a CVE.
Inline:	A device in the active network traffic path. These devices can alter traffic in real-time.
Intrusion Detection System:	A passive system placed in the network architecture to detect and alert on network events based on defined signatures. An IDS does not take action on identified or tagged traffic but can generate alerts

© 2019 The SANS Institute Author retains full rights.

	that are sent to a SIEM or other security response system for mitigation and action.
Intrusion Prevention System:	A system placed within the network path designed to mitigate identified threats based on defined signatures.
Key Terrain, Cyber (KTC):	As defined by JP 2-01.3, key terrain is "Any locality, or area, the seizure or retention of which affords a marked advantage" to an adversary. From a cyber-perspective, key terrain focuses on those assets within an organization which is of operational and strategic significance (Joint Staff, 2014).
Mean Time to Restore (MTTR):	The amount of time that elapses between the start of an outage or impact and the subsequent resolution.
Next Generation Firewall:	Also referred to as an application-aware firewall, these firewalls are capable of analyzing traffic characteristics beyond layer four of the OSI model. They often incorporate IPS-like functionality, thus combining the firewall and IPS solution into a single chassis.
North/South Traffic:	Traffic entering or leaving a given network enclave (host <-> internet, server <-> internet). (Scarfò, 2011)
Passive Device:	A device residing outside of the active network path. These devices typically rely on a network tap or SPAN port to receive the data required for analysis. Passive devices are used for monitoring, analysis, and alerting of traffic, or traffic characteristics, by automated processes or through the actions of a real-time traffic analyst. (Scarfone & Mell, 2007)

© 2019 The SANS Institute Author retains full rights.

Real-time Traffic Analyst (RTA):	An RTA is a network defense operator providing analytics and interpretation of events occurring within a segment of traffic.
Respond:	An action, or the ability to take action, against a detected behavior or characteristic in a traffic segment.
Risk:	The potential for loss, damage, or impact resulting from a given vulnerability or threat vector.
Security Incident and Event Manager (SIEM):	A tool, or suite of tools, designed to aggregate information about security event or incidents. SIEMs range in functionality from log aggregators, to event correlators, to the more advanced security orchestration, automation, and response (SOAR) capability.
Security Orchestration, Automation, and Response (SOAR):	A suite of capabilities, generally associated with a SIEM solution, designed to automate security operations across multiple tools. SOAR capabilities typically include automated mitigation or playbook action execution on a confirmed event.
Signature:	A logical set of conditions within the IPS responsible for categorizing and tagging of traffic that meets the defined criterion. A signature serves as a foundational component in determining an action to be taken against a given flow. **NOTE**: As different vendors refer to signatures using different terms, and for the context of this paper, a signature refers to rules, filters, or definitions as it relates to IPS detection capabilities.

© 2019 The SANS Institute Author retains full rights.

Threat:	Anything capable of exploiting a given vulnerability.
Tuning:	The continuous evaluation of a signature set to achieve an optimal balance between operational capability and overall organizational security posture.
Vulnerability:	A gap or weakness in a given system.

3. Signature Management - Evaluation

3.1. Overview

This section will outline the basic tenants for signature management within an IPS solution. Techniques covered below will address the signature management process from initial evaluation through subsequent employment and continual review, re-evaluation, and expiration. This paper provides a strategy for establishing a continuous signature management policy along with techniques for evaluating an overarching defense-in-depth posture through a multi-layered defense. Please note, this lifecycle process was originally developed for internal use through my work on enterprise-grade IPS appliances and incorporates feedback from various defense department cybersecurity teams.

The following signature lifecycle process will serve as the reference methodology used throughout this paper.

© 2019 The SANS Institute, Author Retains Full Rights

© 2019 The SANS Institute Author retains full rights.

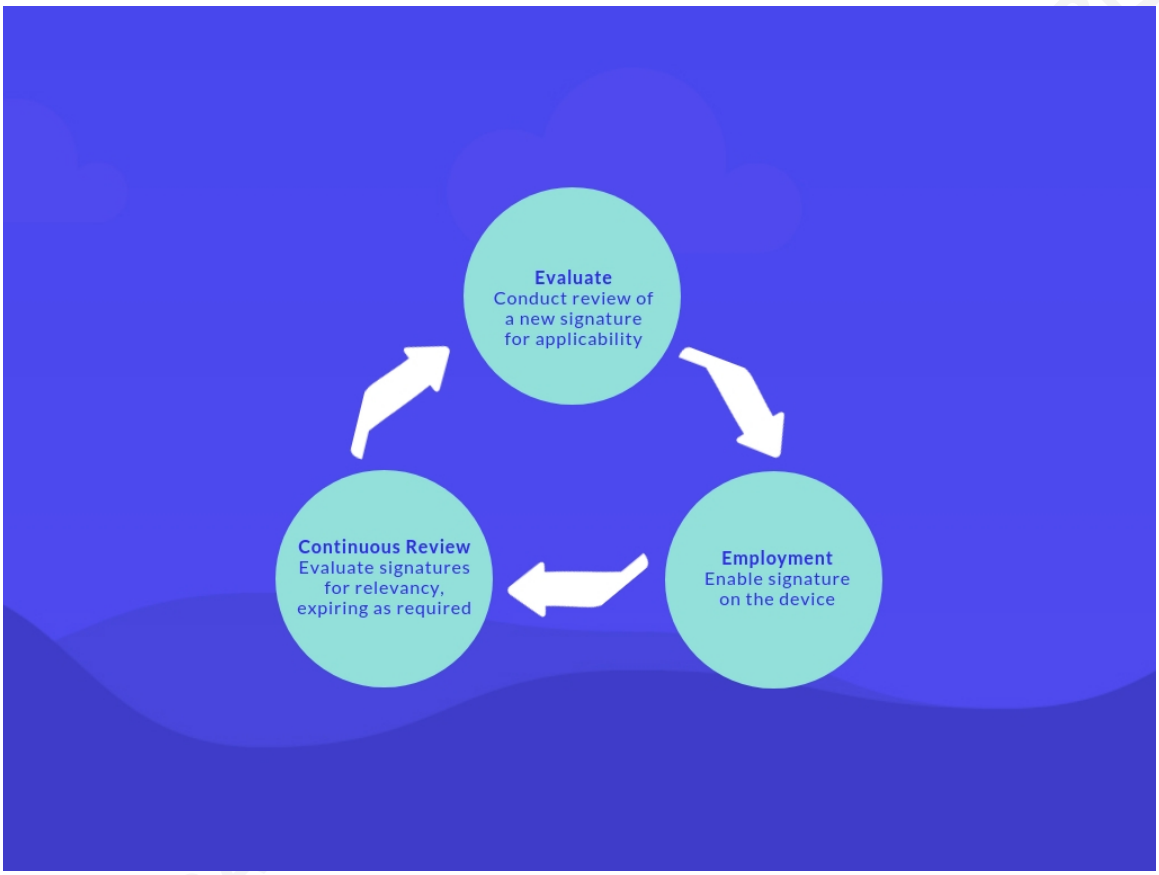

Figure 3.1 – IPS Signature Management Lifecycle

The above lifecycle represents a holistic strategy for the management of signatures within an IPS. As a key component for determining actions taken against matched traffic, proper employment of signatures is critical to network security, stability, and availability. A single improperly tuned or evaluated signature can result in significant impact to IPS device operation and may subsequently impact traffic throughput caused by IPS resource contention. In that light, it is important to highlight that for a properly tuned IPS signature set, the number of signatures employed should have little to no impact on device performance. For this reason, signature evaluation and placement must leverage a holistic defense-in-depth security architecture.

3.2. Traffic Characterization

One of the preliminary steps required before beginning signature and impact assessments is to characterize what normal traffic looks like on a given network segment. What are the top talkers on your network? What volume of encrypted traffic traverses the network that is unable to be inspected? Is there an abundance of specific protocol traffic

© 2019 The SANS Institute Author retains full rights.

which could lead to high evaluation counts for a specific signature? For example, if you have a large amount of SMB transfer traffic on your network, each signature that is employed has the potential to impact device performance for the IPS and network performance of the underlying transport architecture.

Routine characterization of network traffic is a task which must be continuous. Characterization should support the organization's operational tempo and technology lifecycle replacement process. At a minimum, characterization should occur after the introduction of new technologies to the network or upon completion of a change to work center processes. As an example, if the organization is employing a new video conferencing solution, an examination should be conducted to identify the type and volume of traffic produced by this tool. A comparison is subsequently completed against existing signatures employed on a deployed IPS solution to determine the potential for device or service performance impact.

3.3. Evaluation

The first step to the employment of a signature on an IPS is to evaluate the signature from various optics. Upon discovery of a vulnerability, and after an associated signature is developed and released to the community, an analysis must be performed to determine if the IPS is the proper device to host the signature. While most IPS devices are capable of inspecting and reacting to traffic across all layers of the OSI model, the amount of work required to perform these actions on an IPS may be significantly different from a device tailored to specific traffic types. For example, while an IPS can serve as a web filter by restricting access to or blocking specific URLs through a signature, a forward web proxy would be better suited for this action. This placement decision is due, primarily, to processing optimizations and protocol analytic efficiencies present in these devices for processing this type of traffic.

First, a signature must be evaluated to determine the level of risk the threat being mitigated represents to the organization. Presented below is a generally accepted risk determination matrix. Each of these criteria is subjective to an organization and, as such, quantitative and qualitative metrics must be developed to assess impact levels and the likelihood of occurrence.

© 2019 The SANS Institute Author retains full rights.

Risk Assessment Matrix		Likelihood					
		Remote (< 1%)	Very Unlikely (1-10%)	Unlikely (10-25%)	Possible (25-70%)	Likely (70-99%)	Almost Certain (> 99%)
Impact	Catastrophic	6	12	18	24	30	36
	Critical	5	10	15	20	25	30
	Major	4	8	12	16	20	24
	Moderate	3	6	9	12	15	18
	Minor	2	4	6	8	10	12
	Insignificant	1	2	3	4	5	6

Figure 3.2 – Sample Risk Assessment Matrix

Risk management requires an understanding of the overall threat posture within, along with the risk appetite for, an organization. Determining this posture requires several steps to quantitatively and qualitatively analyze the risks surrounding core business assets. While these core assets should serve as the focal point for risk level acceptance determinations, organizations must ensure that lesser known and lower level assets are not overlooked, potentially providing an easy entry point into the enterprise architecture. For a risk analysis as it relates to signature employment, an examination of various aspects of the organization is required to determine the need, and length of time, for employing a given signature:

- **Acceptable risk level** – If the signature employed is to mitigate a current CVE or IAVM, what is the acceptable level of risk that the authorizing official (AO) is willing to accept for any given threat? What protection mechanisms does the AO require once a threat has reached a specific level of mitigation?
- **Patching timeline** – Does a patching timeline exist from the product vendor?
- **Expiration timeline** – At what point can a signature be expired? How does signature expiration tie into the organization's overall CM lifecycle process?
- **Risk applicability** – Is the technology or threat vector present within the organization?
- **Mitigation Method** – Is this threat being mitigated elsewhere in the organization? Are there better technologies to use to mitigate the risk (host-based, network-based, etc.)?
- **Impact assessment** – Will this signature have an adverse impact on network performance?

3.4. Acceptable Risk Level

To determine the need for signature employment on an IPS, the organizational AO must first define the appetite for risk acceptance. Acceptable levels of risk will drive the necessity for protections deployed to address a specific threat. Does a threat require a

© 2019 The SANS Institute Author retains full rights.

95% patch rate before the AO is willing to accept the risk? Is there another protection mechanism or procedure change that can be employed to mitigate this risk to an acceptable level?

3.5. Patching Timeline

For any threat being mitigated through an IPS signature, is there a projected patch that can help define the length of time a mitigation measure needs to be in place? Below is an example of a methodology for identifying, analyzing, employing, and subsequently retiring an IPS signature based on a known patch timeline.

1. An operating system vulnerability with CVE 2018-002 is released on 10/3/2018 with an IPS signature available on 10/6/2018.
2. The signature to protect against this vulnerability is enabled on the IPS on 10/8/2018 after a two-day evaluation period.
3. The vendor announces and subsequently releases a patch on 10/20/2018 to mitigate this vulnerability.
4. The patch is deployed organization-wide to all affected workstations after a three-day test period on 10/23/2018.
5. By 10/30/2018, a 90% patch rate has been achieved. The patching level aligns with the AOs defined criterion for risk acceptance.
6. With the risk level within compliance, the signature is disabled/expired on the IPS and a plan of actions, and milestones (POA&M) employed to cover the remaining 10% of workstations.
    a. Optionally, a matching signature may be employed on an IDS to detect any residual traffic matching the original IPS signature.

As patching timelines will vary from product to product and vendor to vendor, each threat or vulnerability must be evaluated on a case-by-case basis.

3.6. Signature Expiration Timeline

At what point can a signature be expired or disabled on an IPS? As mentioned with the patching timeline example, expiration or disabling of a signature should tie into the overall risk management, patch management, and lifecycle management processes of the organization. When a product is removed from an organization's cyber inventory either

© 2019 The SANS Institute Author retains full rights.

© 2019 The SANS Institute, Author Retains Full Rights

through upgrade attrition or product removal, the associated protection techniques (signatures, products, rules) should be incorporated into this lifecycle process for analysis.

As with the previous CVE example, once a specific risk mitigation level is reached, organization policy could direct the disabling of a signature with an understanding that further actions will be taken to continue to reduce the number of unpatched systems. From another optic, when a product is removed from within the organization's boundaries, an analysis should be performed to determine if the associated signatures can be disabled.

3.7. Signature Applicability

In line with the patching timeline discussed above, each signature must be evaluated from a standpoint of applicability to the organization. Organizational change and configuration management solutions can assist in determining whether a specific signature or protection is required.

While an argument is valid that any zero-day threat should have protective measures employed to mitigate the risk introduced to an environment, if the exploitable product is not present within the digital boundaries of the organization, is protection required? Having a firm understanding of the products, applications, or appliances deployed throughout an organization can reduce the number of unneeded signatures employed on an IPS. As an example, are signatures designed to detect threats pertinent to Android mobile phones necessary on a network without such devices? Should signatures be enabled to prevent a zero-day in a web-based application not installed within the organization's boundaries? This level of understanding will result in a reduction of processing overhead on installed security appliances and organization endpoints.

3.8. Threat Mitigation Method

The next piece of the puzzle to consider before implementing an IPS signature is whether or not a more suitable solution exists to mitigate a given threat. While most IPS solutions protect traffic at all layers of the OSI model, they may not be the most efficient system to achieve a desired level of protection. For example, on a network with a dedicated proxy, consolidating domain or URL blocking rules to this proxy can ease troubleshooting efforts, simplify configuration and change management, and alleviate the

© 2019 The SANS Institute Author retains full rights.

potential load on other inline components. Different technologies exist to protect against threats for a specific network, transport, and application layer protocols. These technologies should, ideally, be employed utilizing a layered methodology to ensure that a threat missed by one system is caught elsewhere in the inspection chain.

3.9. Signature Impact Assessment

When employing a signature, an analysis of the impact the signature could have on an IPS is required. This analysis should look into the following aspects of traffic and device performance to determine the potential for impact:

- The expected evaluation time for a single hit against the signature
- The expected volume of traffic that could potentially trigger the signature for inspection
- The likelihood of false positives against the signature
- The tuning characteristics for the signature that could increase or decrease signature hit counts

These items each tie into the level of impact to be expected from a given signature (Schaelicke et al., 2003). The evaluation time can relate to potential delays for traffic reaching its destination within the destination network. The expected traffic volume for a given signature and the likelihood of false positives can lead to resource contention issues within the IPS. Tuning characteristics and signature heuristics can significantly increase or decrease hit counts, potentially leading to resource availability issues or false negatives. A single, improperly tuned signature can have a more drastic impact on device and network performance than a device running thousands of properly tuned signatures.

An example of a situation where a single signature could impact the rest of the device is a signature designed to inspect SSL traffic for malicious or malformed certificates. When a signature of this nature is enabled on a network with a high volume of SSL traffic, this single signature may significantly increase load by causing excessive evaluation per SSL flow. Excessive traffic evaluation may lead to inspection queueing as processor I/O operations increase. This impact may manifest as latency for network applications or services which would further hinder network performance across the organization.

© 2019 The SANS Institute Author retains full rights.

3.10. Web URL Filtering Example

The next example expands the URL-blocking situation mentioned earlier in this paper. Many IPS vendors provide signatures to detect and respond to specific web-based traffic exploits and threats. Likewise, most IPS vendors also offer methods for restricting access to known bad, or organizationally-prohibited, URLs and web resources. There are multiple challenges with implementing these blocks on an IPS. The first, and perhaps most important, challenge is that the IPS is likely less optimized compared to a forward web proxy to inspect and take specific actions against web-based traffic. This lack of optimizations is not to say that there are no optimizations within the IPS for this type of traffic, but that a proxy has been purpose-built for processing of web-based protocols (HTTP, HTTPS, etc.). These optimizations lead to processing efficiencies not present within the IPS which, in turn, reduces the resources required to provide protections similar to the IPS-based signature.

The next aspect to consider is the implementation and management of URLs contained within a given IPS signature. When utilizing an IPS-level block, the level of administrative overhead necessary to maintain IPS signature configurations must be understood. Another aspect to consider is the amount of effort required to maintain and manage existing URLs programmed into a signature. This management overhead includes updating and distributing signatures across the entire IPS architecture, in the case of distributed enterprises.

Exacerbation of this overhead occurs when utilizing multiple IPS signatures to perform similar functions. A change management solution is required to track the functions being performed by individual signatures. Additionally, minimizing complexity for the transport and inspection architecture can reduce MTTR for outages by keeping protection mechanisms on their associated management solutions. This restoral time can be hindered by having multiple systems, or multiple signatures, in place providing similar functionality.

4. Signature Management - Employment

After completion of the above analysis, organizational processes and procedures should dictate whether or not the signature is employed. Upon employment, an impact assessment period should be established to determine if any unforeseen performance

© 2019 The SANS Institute Author retains full rights.

impacts occur. While testing and evaluation of signatures can eliminate potential network and performance impacts, Murphy's law generally creeps up once a new capability moves into production. For this reason, each signature or set of signatures should be evaluated for any unwanted second or third order effects of implementation following a standard change management model.

For signatures approved for deployment, the organization should execute the previously mentioned change period. The change management period will standardize the time frame during which new signatures are employed and no longer required signatures are disabled. This change period will allow network technicians and RTAs to assess performance using a predictable schedule while also minimizing the chance for impact on organizational operations. Through this method, RTAs can rapidly identify any impact introduced by new signatures.

By leveraging an evaluation window after a change to an IPS, the MTTR for an outage is minimized by adhering to change management processes, allowing for a rapid roll-back of an impact causing change. As with any information technology change, maintenance windows should be established that address changes to specific systems or technologies. It would complicate the troubleshooting process, for example, to implement a weekly patch cycle at the same time you are tuning the IPS.

5. Signature Management – Continuous Review

As with any policy or procedure, continuous review and improvement should remain a part of any signature management strategy. The same steps recommended to evaluate signatures should also be used to perform a continuous review of enabled signatures. Codifying these steps into a formal review process will ensure computing resources are not wasted defending against threats not present, or reduced to acceptable levels, within an organization's architecture.

Through proper review, new and old signatures may be appropriately maintained, ensuring a demonstrable security posture for the organization. This process also supports the lifecycle of other technologies within an enterprise. By integrating signature management into the overall enterprise lifecycle replacement process, the organization can ensure that network traffic is not being unnecessarily inspected for threats against technologies no longer present in the business's portfolio.

Josh Levine, josh@cantreach.me

© 2019 The SANS Institute Author retains full rights.

6. Signature Tuning

6.1. Overview

A key component of any signature management strategy is proper signature tuning. As mentioned previously, there are multiple facets to a signature that can impact inspection times, device operations, and overall network performance (Erlanger, 2004). These performance metrics can be improved by focusing on specific tuning criteria outside of the actual enabling/implementation of each signature.

6.2. Heuristics

Signatures released across various IPS platforms support a heuristic-based classification of traffic. These heuristics allow for characterization based on machine-learning, artificial intelligence, or other methods to identify traffic that may fall slightly outside of a signature's specified criteria. This capability allows for signatures to be applied dynamically, potentially lessening the need for subsequent signature updates as the detection characteristics change.

As with any dynamic detection capability, it is important to exercise caution when tuning the heuristic thresholds. An example of such tuning would be for an email inspection signature designed to block messages potentially flagged as spam. While this signature may include a default setting only to inspect traffic where there is a seventy percent chance of the email being spam, tuning this setting to a lower value could have an impact to IPS and email service performance. For example, setting the heuristic threshold to twenty percent requires a very low likelihood of a message being spam before the execution of an inspection or blocking action. These actions may increase device utilization. The second and third order effects of this could be delays in message delivery with the potential for blocking of legitimate message traffic.

For this reason, tuning any dynamic detection capabilities of a signature should be conducted carefully with appropriate service owners and network administrators notified ahead of time. Proper coordination will allow the teams to maintain awareness for adverse impacts should the tuning have unexpected results.

Josh Levine, josh@cantreach.me

© 2019 The SANS Institute Author retains full rights.

6.3. Trusted Traffic and Traffic Bypass Rules

While the idea of purposely excluding traffic of any type from inspection seems counterintuitive to the basic tenants of security, there may be situations which warrant some form of traffic bypass. These exclusions should be the exception and not the norm but can prove helpful in certain situations.

As shown above, the implementation of a signature can have a direct correlation to impact on an IPS device or the underlying transport network as a whole. For this reason, trusted traffic rules are provided in most IPS solutions as a means to bypass traffic which is high in volume, unable to be inspected (think encryption), known to be safe or low risk from exploitation, or protected through other means. This type of bypass can result in a lower potential for impact to device and network performance but can also present a risk regarding a lowered security posture.

When looking at high volume traffic, it is important to take into account the directionality of traffic and whether the flows typically occur between the same, or similar, groups of hosts. As an example, for software distribution and management systems such as Microsoft's System Center Configuration Manager, examining the deployment architecture for primary management servers and secondary distribution nodes can identify points in the network where trust rules may be established. These rules would allow traffic to bypass the inspection engine, instead relying on host-based mechanisms to detect threats and vulnerabilities, at the IPS-level. Since these systems use a defined architecture with a standard set of ports and protocols, it is simplistic to limit the inspection of this high-volume traffic using standard IPS traffic management and trust rules. The simplicity of exempting routine traffic from inspection by an IPS, offloading this inspection to another solution within the enterprise, highlights the need to maintain an understanding of how deployed applications are architected and communicate.

Another example of traffic where a trust filter may be beneficial in reducing the load to an IPS is for backup or replication systems such as NetApp's SnapMirror solution. Due to the high volumes associated with this type of traffic, trusting traffic between the source and destination storage arrays can alleviate the need to inspect traffic over known SnapMirror ports between trusted endpoints. Please note, this trust rule should not trust

© 2019 The SANS Institute Author retains full rights.

all traffic leaving one or both hosts and should be implemented in as restrictive a method as possible (source<->destination port/IP vice source<->destination IP).

A third example for trusted traffic focuses on systems that generate large amounts of traffic on a routine basis. In this particular case, the culprit is a vulnerability scanner. As these systems are designed to detect applications or other systems on a network that may be susceptible to vulnerabilities, the traffic they generate may contain characteristics of known malicious activity to elicit a response from the target. In line with the evaluation examples provided in this paper, this type of traffic presents a new opportunity for additional appliance load. Since this traffic will likely mirror activity targeted by threat signatures, alerts and subsequent blocking actions may be taken by the IPS to prevent the traffic from reaching the target.

For this reason, traffic egressing the vulnerability scanner may present a possible trust rule violation scenario. Additional protections should be employed to ensure that inspection occurs for traffic originating outside of known scanning windows. This methodology will account for a scenario where the vulnerability scanning system is compromised through lateral movement or similar vector and binds to an excepted IP/port pairing.

Traffic bypassed from inspection should be documented and approved by the organizational AO. An analysis of protection mechanisms available within the organization should be used to determine if any risks introduced through a trusted traffic bypass are mitigated elsewhere in the architecture. As with any robust security architecture, these layered defenses can assist in justifying a bypass to leadership.

IPS solutions differ in their implementations for trusted traffic rules and how they operate. Some appliances offer both capabilities where a trusted traffic rule is used to process traffic through the device but not actively inspect it in software. Conversely, a bypass rule can be implemented at the IPS hardware level to identify ingress traffic based on a given port number and immediately forwarding the matching traffic out of the IPS egress interface.

The challenge with trusted traffic rules that are port-based and not based on protocol-level header information is that this potentially opens a hole into the network for traffic masquerading on a known, trusted service port. For example, if a backdoor Trojan finds

© 2019 The SANS Institute Author retains full rights.

its way into the network enclave and uses probing techniques to detect allowed ingress/egress ports, it could bind to this port thus allowing for uninspected traversal into and out of the network. As such, trust and bypass rules should be used with caution and only when other detection mechanisms exist to identify malicious or anomalous traffic. Additionally, any traffic that is explicitly trusted or bypassed at an IPS should be inspected elsewhere in the architecture.

6.4. Signature Tuning – SSL Trust Rule Example

A real-world example of a trusted traffic scenario is the inspection of SSL traffic passing through an IPS. Using a defense-in-depth model, leveraging host-based SSL scanning technologies such as client-side certificate validators or host-level intrusion prevention or firewall systems, resource load on the network-based IPS can be reduced. This reduction is achieved through the process of bypassing SSL traffic at the IPS for TCP port 443. In an environment that does not possess an inline SSL decryption capability, this traffic can present a significant load on an IPS with little ROI based on available signatures for SSL traffic inspection. By its very nature, the payload for any SSL traffic is encrypted. This encryption limits the abilities of an IPS to provide real-time protection and response outside of certificate or certificate authority validation and SSL traffic flow characterization (volume, duration of a conversation, packet size, etc.).

The inspection of this traffic can instead be shifted from the network layer to the host layer, relying on host CPU cycles to perform this decryption and inspection. This methodology will enable other host-level protection techniques through products such as data loss prevention tools, host-based firewalls, and host-based data execution protection measures after the traffic is unencrypted at the host. This shift in traffic analysis can provide an increased security posture while reducing the operational load to installed network-based security appliances.

Using the above example, let's expand it by using a sample signature designed to evaluate an SSL traffic flow for certificates issued by a known compromised certificate authority. For this example, please note that calculation timeframes are not specific or intended to be realistically accurate but instead are highlighted to serve as a reference for the level of resource impact which may occur on an IPS.

Scenario characteristics:

© 2019 The SANS Institute Author retains full rights.

s = Signature evaluation time (per traffic flow): .5ms

f = Average number of SSL traffic flows per minute on this segment: 2500

T = Total processing time required for this signature

Example calculation formula for this signature: *T = (s * f) / 1000:*

*T = (.5 * 2500) / 1000*

T = 1250ms / 1000

T = 1.25s of computing time per minute per signature

For the above example, this single signature requires 1.25 seconds out of every minute to evaluate 2500 SSL encrypted traffic flows for a single known malicious certificate. Expand this evaluation to include all signatures currently enabled on the IPS profile, and it is easy to see how quickly processing resources can end up in a contentious state (Ethala, Seshadri, Renganathan, & Saravanan, 2013).

It is for this reason that determining the ROI expected for a given signature be evaluated, and appropriate device selected, for protection or detection of this specific threat. If placement of a signature on a passive device is acceptable, a reduction in load to the IPS occurs while detection is still possible within the organizational security construct. (Scarfone & Mell, 2007).

7. Defense-in-Depth

7.1. Overview

Defense-in-depth is one of the foundational tenants of modern cybersecurity. The basic idea is to layer varying security technologies on top of each other to protect against threat vectors utilizing different protection methods. These protections should span the entire spectrum of security ranging from all layers of the OSI model to include the human layer. Another tenant behind defense-in-depth is that these layered protections will serve as back-stops should one layer fail to detect or prevent a threat. For this reason, it is important to understand the security architecture employed within an organization. A clear understanding will help ensure the primary defense for specific threats is conducted using the right technology with these back-stops serving to fill the gap.

© 2019 The SANS Institute Author retains full rights.

7.2. North/South vs. East/West Protection

When talking about network traffic inspection, there are two types of traffic flows that are of primary importance. The first is for traffic entering or leaving a network enclave, commonly referred to as north/south traffic, while the second type occurs between hosts or servers on a given network, referred to as east/west traffic. This type of traffic takes on increased importance for analysis when it comes to setting up internal inspection trust filters, discussed earlier in this paper. This traffic includes movement between individual hosts or pre-defined communication patterns between a server and a set of hosts on a network.

When examining the methods for traffic inspection, it is important to understand what a given signature is designed to detect or prevent. For example, is it necessary to inspect SMB traffic egressing the network? It is a commonly accepted practice that blocking of SMB traffic occurs at the border firewall. For this reason, there shouldn't be a necessity to provide SMB traffic inspection at the boundary (US-CERT, 2017). SMB inspection takes on greater importance regarding inspection within the network enclave where the desire to detect lateral movement exists. While it may not hurt to inspect SMB traffic going north/south, the primary concern for this type of traffic exists internally.

For this reason, it is important to add directionality of inspection to an IPS signature evaluation criteria. Directionality is another reason why placement of an IPS takes on increased importance as the organization determines which traffic it is interested in inspecting. For any defense-in-depth posture, inspection of east/west and north/south can play a role depending on the definition of the organization's KTC.

7.3. Host-Based vs. Network-Based Protection

As with other defense-in-depth scenarios, network and host-based detection mechanisms vary with regards to traffic processing and inspection capabilities, functionality, and optimization. Network-based capabilities generally focus on detecting items on the wire before reaching a specific endpoint. Host-based detection and response capabilities focus on traffic once it has reached its destination and, generally speaking, have access to a larger array of data as the target applications de-encapsulate, decrypt, and process the received traffic.

© 2019 The SANS Institute Author retains full rights.

For situations where requirements include things like playback of traffic conversations, analysis of traffic conversation statistics (flow data), or where protection is necessary before the traffic reaches a target host (browser and web-based vulnerabilities, for example), network-based detection and prevention are preferential. When a requirement exists for more in-depth inspection for application-level data, host-based detection and prevention is the preferred choice. After all, some data is only available once it has been decrypted and processed by the receiving endpoint or application.

Analysis of the type of information and which subsequent detection or response action is required within the organization will shape the placement of signatures in a defense-in-depth security architecture. Based on previous recommendations, the IPS may not be the best place to perform the inspection of specific traffic or conversation types. The IPS can be used to offset or augment some of these inspections, but the primary method for detection and response may be another appliance entirely.

Other detection and protection mechanisms in a defense-in-depth security architecture include technologies such as forward and reverse web proxies, layer 3/4 firewalls, next-generation or application-aware firewalls, inline multi-vendor antivirus scanning tools, file sandboxing solutions, or passive devices such as an IDS or SIEM. Each of these technologies provides a different inspection layer for traffic traversing a network and associated endpoints.

7.4. Detect vs. Respond

The discussion around host or network-based actions brings us to the next topic of performing active response or resorting to passive detection and alerting. You're your organization have a requirement to respond to a threat identified on the network or is detection and alerting enough? Ultimately, this decision comes down to the risk posture for the organization and what the AO is willing to accept.

Part of any good defense-in-depth strategy includes passive detection capabilities. These capabilities typically rely on a network tap or SPAN port that feeds a passive IDS device. These systems can serve as a fallback for an inline IPS and provide alerting to personnel on threats detected within, or transiting into/out of, the network boundary. Since the primary difference between an IDS and an IPS is the ability to take responsive action, the signature sets should be reasonably similar to allow for mirroring of signature

© 2019 The SANS Institute Author retains full rights.

sets between active and passive devices. This technique allows for a higher number of passive IDS devices to be emplaced throughout the network to detect traffic that may have been passed as a false negative by the inline device.

While an active response is an excellent option for organizations that have the resources and staffing to tune these devices, passive detection can serve to alert on potentially malicious traffic. With many free and widely supported products available for detection, it is possible to deploy these solutions in a manner that provides greater network visibility without the overhead/risk of running a full-blown intrusion prevention system.

7.5. IPS Logging and Debug Actions

Another option to look at for identifying the impact of a given signature is to utilize built-in alerting and trace alerting capabilities of the IPS/IDS. Several solutions allow signatures to be set-up in logging mode with an optional debug option. While logging mode is a decent method to use when evaluating a signature's impact, it is important to understand that logging and debugging have the potential to negatively impact device performance, leading to second and third order network performance impacts (Newman, 2006). For this paper, the following terms will be used to highlight specific logging scenarios:

- Allow – Traffic is permitted through the device
- Allow + Alert – Traffic is permitted through the device and an alert is generated.
- Allow + Alert + Debug – Same as Allow + Alert but additional debug information is captured related to the specific conversation (bytes in/out, packet capture, flow records).
- Deny – Traffic is blocked at the device
- Deny + Alert – Traffic is blocked at the device, and an alert is generated.
- Deny + Alert + Debug – Same as Deny + Alert but additional debug information is captured related to the specific conversation (bytes in/out, packet capture, flow records).

The Allow and Deny options are straight forward in their function in that the device either permits or blocks a specific traffic flow. For the alert options, one or more logging events take place related to the traffic flow. These actions may include logging to the

© 2019 The SANS Institute Author retains full rights.

internal device log management system, or the generation of an alert sent to an external SIEM. The debug option takes alerting one step further and produces additional artifacts to assist in troubleshooting situations where traffic might not pass through the device as expected. These artifacts can include flow data for a given traffic conversation, a full packet capture, device-specific debug outputs and other miscellaneous data.

Each of these options fits different scenarios when it comes to enabling or disabling a given IPS signature. The challenge with utilizing these options is ensuring they do not negatively impact the device and the underlying network architecture. Let's take the previously mentioned example involving high volumes of SSL traffic and apply it to a signature designed to inspect the header information of an SSL conversation.
By enabling a rule in an Allow + Alert or Allow + Alert + Debug configuration, a significant load could be introduced to the device for each SSL traffic flow now generating an alert or debug activity. The unintended effect of such a change could lead to resource contention within the IPS. This impact could manifest as latency, delayed application flows, or dropped traffic as the device struggles to keep up with new alerting and debugging tasks (Markatos, Papadogiannakis & Polychronakis, 2010). The outcome of such an impact could result in a device operating in failback mode. This likelihood increases should technicians be unaware of the alert or debug setting enabled on a given signature or the number of evaluations or events generated by the signature. Understanding debug options highlights the need for a well-defined and tracked change management process, as it relates to signature evaluation and management, and is paramount to a successful implementation of signatures using this method.

8. Conclusion

This paper was intended to provide a vendor-agnostic view on IPS signature management theory. I hope this paper illuminates some areas to consider when developing strategies for your organizations. Each of the concepts presented here is designed to serve as a base guideline as a single methodology that fits all organizations does not exist. How one organization develops and employs signature management techniques can vary significantly with another.

By establishing an effective and focused signature management strategy, the overall security posture of an organization can be greatly enhanced using a methodology that supports a balance between functional operations and security. Continuous review and

© 2019 The SANS Institute Author retains full rights.

refinement of this strategy will ensure signatures employed relate directly to an organization's KTC.

© 2019 The SANS Institute, Author Retains Full Rights

Josh Levine, josh@cantreach.me

© 2019 The SANS Institute Author retains full rights.

9. References

Barnum, S., Gegick, M., & Michael, C. (2005, September). Defense in Depth. Retrieved from https://www.us-cert.gov/bsi/articles/knowledge/principles/defense-in-depth

Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. ACM computing surveys (CSUR), 41(3), 15.

Erlanger, L. (2004). Intrusion Detection Needs a Dose of Prevention. *InfoWorld, 26*(11), 42-45.

Ethala, K., Seshadri, R., Renganathan, N. G., & Saravanan, M. S. (2013). A role of intrusion detection system for wireless LAN using various schemes and related issues. 10, 979-985. doi:10.3844/ajassp.2013.979.985

Joint Intelligence Preparation of the Operational Environment. (2014). Joint Staff Retrieved from https://fas.org/irp/doddir/dod/jp2-01-3.pdf

Markatos, E., Papadogiannakis, A., & Polychronakis, M. (2010). Improving the accuracy of network intrusion detection systems under load using selective packet discarding. Paper presented at the Third European Workshop on System Security, Paris, France.

Mitre. (2018). Common Vulnerabilities and Exposures. Retrieved from https://cve.mitre.org/

Newman, D. (2006). IPS: Slow down for safety. *Network World, 23*(35), 38-38,42,44,46,48.

Scarfone, K., & Mell, P. (2007). NIST Special Publication 800-94, Guide to Intrusion Detection and Prevention Systems (IDPS). Retrieved from https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-94.pdf

Scarfò, A. (2011, June). The evolution of data center networking technologies. In Data Compression, Communications and Processing (CCP), 2011 First International Conference on (pp. 172-176). IEEE.

Schaelicke, L., Slabach, T., Moore, B., & Freeland, C. (2003). Characterizing the Performance of Network Intrusion Detection Sensors. Paper presented at the International Workshop on Recent Advances in Intrusion Detection.

US-CERT. (2017, March 16). SMB Security Best Practices. Retrieved from https://www.us-cert.gov/ncas/current-activity/2017/01/16/SMB-Security-Best-Practices

© 2019 The SANS Institute Author retains full rights.