Evaluating Enterprise IPS: Seven Key Requirements

Joel Snyder / Opus One

Abstract/Executive Summary:

Network Intrusion Prevention Systems (IPS) can be effective tools in the arsenal of the security practitioner—or they can be expensive failures. To help ensure the former, an enterprise IPS deployment should meet seven key requirements. Security managers looking to install IPS in their networks will be able to use evaluation criteria in this white paper as a way to help differentiate products and identify those most suited to an enterprise deployment.

The seven requirements for enterprise-class IPS are:

1. Availability of a broad hardware range with offerings ranging from branch-office T1-class devices to enterprise multi-gigabit systems;
2. A detection engine incorporating multiple technologies, including signature matching and protocol anomaly and behavior anomaly detection;
3. An ability to provide a "big picture" context of information about the network and its devices to qualify security information as well as speed analysis;
4. IDS-like features to detect less critical threats, perform behavioral analysis and facilitate security forensics toolkits;
5. Centralized management and alerting capabilities across multiple devices;

6. An analysis toolkit that efficiently drives policy refinement; and,
7. Network security tools, especially firewalling, to simplify deployment

Introduction:

Intrusion Prevention Systems (IPSes) and Intrusion Detection Systems (IDSes) have had a mixed history with the security community. Although adored by the security purist as invaluable weapons in the defense of the network because of their intense visibility into threats and attacks, practitioners in the enterprise have shied away due to a host of faults, ranging from performance slowdowns and false positives to a general lack of manageability.

These issues have not deterred security vendors from developing, refining, and improving their products, however.

At the same time, external pressures from compliance regimes like the Payment Card Industry and HIPAA regulators are requiring enterprises to have strong security verification programs in place. The recent raft of reported data breaches, along with their enormous mitigation costs, has also given network security architects reason to put IPS within the network. In short, the cost and aggravation of doing IPS is quickly being offset by the cost and aggravation of **not** doing IPS.

That said, it's time for a second look at IPS products. This white paper is primarily focused on enterprise requirements for IPS. However, we cannot consider or discuss Intrusion Prevention without also considering Intrusion Detection.

The primary goal of an IPS is to block attacks. This means that IPSes generally have a small set of very tightly defined signatures with a very high expectation for accuracy. False positives or performance delays in an IPS can affect production network traffic, so these systems are engineered and programmed to avoid both problems.

In contrast, IDSes are designed to detect security problems, which may include everything from bona fide attacks to mis-configured applications and firewalls. While an IPS is a threat mitigation tool, an IDS is more like a "protocol analyzer" for the security analyst. Because of the natural overlap in detection technology at the core of both IPS and IDS, it is common to find IPS products that include some IDS features.

Enterprise network managers considering IPS and IDS technologies need to identify products that are suitable for large-scale deployment. While it's tempting to focus on the detection engine inside of the product, there is more commonality than there are differences across the IPS/IDS market when it comes to detection engines.

Instead, the differentiation across products lie in the host of other features, ranging from performance to manageability to forensics, that really separate out enterprise-class products from those aimed at smaller networks or less sophisticated security environments.

In this white paper, we identify seven key requirements for an enterprise-class IPS. These requirements can be used as evaluation criteria check list to help identify products most likely to succeed in an enterprise network deployment.

1) Performance and Range of Hardware

When evaluating IPS devices, one of the key differentiators is speed with which it can absorb and process network traffic. The requirement for an IPS to not slow down the network is so obvious that it's almost redundant to state it. When called upon to protect gigabit-speed links, the IPS must neither become a bottleneck nor introduce more than a few milliseconds of latency.[1]

IPS vendors certainly find the high end of high-speed devices the most attractive. It not only gives developers interesting and exciting software and hardware challenges, but also makes sales teams happy, with high per-unit costs. However, an enterprise-class IPS must be able to scale up to gigabit (or faster) speeds with minimum latency as well as be able to scale down to branch office size and speed at reasonable cost and in a suitable form factor—the latter being a less exciting proposition. An enterprise IPS must have low-speed hardware and software as good as its high-end, because most deployments will include both somewhere in the network.

Scaling up creates different requirements than scaling down. Scaling up starts with performance, the ability to handle the gigabit speeds of the enterprise. That said, just as important as raw speed is the ability to exhibit graceful and controlled behavior when problems arise. For example, enterprise-class IPS must include "fail-open" hardware so that a system failure does not interrupt network flows.[2]

[1] This requirement also exists for IDS, although since the IDS is outside of network traffic flows, performance problems in an IDS simply affects is accuracy and ability to discover problems, rather than taking down the entire network.

[2] The term "fail-open" can be confusing. If an IPS is a device that sits between two other network devices, the idea is that during a failure, the connection is opened so that all traffic flows unimpeded between the devices; the IPS becomes a "patch cable" rather than an active element. It is considered safe for an IPS to "fail-open," while other security devices such as firewalls always "fail-closed:" they stop all connections rather than risk letting inappropriate traffic through. The confusion of the term comes because in electrical engineering, an open circuit is one that passes no current, while a closed circuit lets current flow—the conceptual opposite of what is happening here. Whichever term you prefer, the idea is the same. When an IPS fails, it should pass all traffic rather than take down your network.

And where there is fail-open hardware, there needs to be "fail-open software." When the load on the IPS goes above sustainable levels, the IPS needs to gracefully degrade its inspection process so that it does not add excessive latency or begin to lose packets. When evaluating this feature, look for sufficient administrative controls. You should be able to control alerting thresholds and other parameters (such as a maximum alerts-per-second or packets-per-second levels) so that when things begin to go wrong, you'll immediately know it and be able to dictate the system's behavior to remain within your SLA and keep the network running smoothly.

Predictable performance and fail-open capabilities aren't the only requirements for high-end hardware. Enterprise deployments may also call for interface flexibility (such as including fiber and copper interfaces on the same system), high-availability design features, and scalability without a "forklift upgrade."

For successful enterprise deployment, the IPS needs to scale down as well. Offering a scaled down IPS suitable for branch offices or remote sites means more than a vendor simply selling a baby box. Scaling down means that everything should be specifically designed to fit the branch office, including the system price. The hardware has to be small, while noise, power requirements, and radiated heat all have to be tightly controlled for office-environment deployment. Integration between IPS and other branch-office security devices, such as UTM firewalls, is important when scaling down as well.

Scaling down at the branch, paradoxically, requires more advanced centralized management. When a hundred or even a thousand small devices are being deployed, this calls for great sophistication in the design of the management console. Features such as compression of alerts that reduce bandwidth consumed, deployment wizards and grouping of devices to make it possible to manage many elements, delegated and hierarchical management capabilities, and automated updating of software and policy parameters are all important to properly scaling down.

2) Wide Range of Detection Engine Technologies

Every IPS includes an engine that handles signature-based intrusion detection inside—that is a given within this industry. While much of the theoretical research on IPS has focused on making the detection engine as strong as possible, the enterprise reality is that differences in signature-based detection engines are almost undetectable and thus do not strongly differentiate products. For example, more than a dozen commercial products are all built upon the open source Snort detection engine or a close derivation of it.

It is for this lack of differentiation–coupled with the fact that simple signature matching is not sufficient to catch the security problems network managers really need to know about—that sets the requirement that an IPS must also employ other means of detecting malicious network behavior.

The reality for most enterprise networks is that an IPS is not going to block many traditional attacks coming from the Internet, because most network managers have done a good job installing and configuring firewalls and keeping their Internet-accessible servers patched. The IPS won't see the attacks because the firewall blocks them, and those layer 7 attacks that get through an intentional hole in the firewall will often be attacking servers that are not vulnerable. The problems that IPSes are going to identify and block are much more subtle and sophisticated, ranging from mis-configured servers and applications to curious insiders and internal assets that have become infected so that they've joined "bot" networks. Signatures, alone, won't help address any of these problems.

Three key features that go beyond standard signatures are: a top-notch security vulnerability research team, layer 7 protocol visibility and knowledge, and network behavior analysis.

Any top-notch team maintaining a good IPS detection engine should be visible in the security community. This doesn't mean a spate of press releases, but a continued and obvious presence both at conferences and in the on-line world where most Internet security is discussed. The team must demonstrate that its members are not just adapting other people's discovered vulnerabilities into their own engine, but are also actively searching for problems and are—at least occasionally—first "out the door" with a new zero-day discovery. This team presence has to translate into action, with measured and continued updates of both the engine and detection signatures. IPSes are not like anti-virus tools; they don't need hourly updates, especially if they've been written properly in the first place. But there will always be new threats, and the team behind the engine needs to be on top of these threats in a timely fashion, in days, not weeks.

Protocol anomaly detection and visibility is another key requirement for enterprise-class IPS. The IPS needs to understand not just layers 3 and 4 (such as IP and TCP/UDP), but also must be able to identify anomalies all the way up to layer 7: deep inspection into this protocol layer in order to identify illegal transactions and upended protocol states, buffer overflows, and other suspicious activities. Clearly, the most vulnerable protocol today is HTTP, but enterprise-class IPS devices need to know about as many layer 7 application protocols as are running on your network, such as DNS and

FTP, mail protocols such as SMTP and IMAP, remote access protocol such as RDP, SSH, and VNC, and even internal-use-only protocols such as Microsoft's email (MAPI), file sharing (CIFS), and RPC protocols, RADIUS, SNMP, TFTP, and dynamic routing protocols.

Network behavioral analysis is another area where IPS detection technology should be extended for enterprise-class networks. Behavior analysis is the science of understanding all traffic flows across the network to identify a common baseline, including factors such as top talkers and listeners, sources and sinks for traffic, general flow direction, as well as the specific TCP or UDP ports in use. Once this baseline is established, the IPS can identify deviations from the baseline as potential security problems. Network behavioral analysis is usually not part of a standalone IPS sensor, because the risk of a false positive with behavioral analysis is dramatically higher than with other types of detection. For example, a seldom-used server may not show up on the baseline, or a slight change in network topology or a disaster recovery event may cause very different network behaviors that are still perfectly legal. Instead, enterprise-class IPS systems should have behavioral analysis as part of their management systems—more of a detection (IDS) function, than a prevention (IPS) function—with tools to establish baselines, and to alert or even block traffic when baseline deviations are detected.

3) Big Picture Context

When analyzing information coming from an IPS, most of an administrator's time is spent scrolling through alerts answering the question "do I care?" (or, less politely, "so what?") The money spent on IPS products pales in comparison to the time it takes for a security staff to properly understand and react to the information coming out of the management system. In the enterprise, an IPS must provide more contextual information about the network, the users, and the systems residing on the network. The IPS must understand which systems are critical, which systems are vulnerable, which users are staff and which are guests, what the network topology is, and how any single alert relates to the user and system it affects.

At the most basic level of functionality, IPS sensors will transmit alerts about dropped packets and broken connections back to a central management system. An analyst reviewing these alerts will quickly become overwhelmed with hundreds, thousands, or even hundreds of thousands of one-line messages about problems the IPS is blocking. Despite the difficulty of this task, the logs from the IPS must be reviewed to identify security issues that have to be addressed or to catch false positives and change IPS policies. While IPSes are not as chatty or error-prone as their IDS cousins, the IPS manager cannot simply use a fire-and-forget strategy in an enterprise network by never reviewing alert information.

In an enterprise environment, this alert review and analysis process absolutely requires more than a long list of messages. The alerts must be put into context which places the affected system within the complex topology of an enterprise network and helps the analyst understand where the "attacking" and "attacked" systems are located on the network. The alert must also have sufficient evidence (such as the packets that make up the problem) attached to it so that the security analyst can decide whether it represents a false positive (incorrectly raised alert), a false alarm (an alert that really isn't a problem in this context), or a true security issue.

To make the best use of the security analyst's time, alerts from the IPS also need to be supplemented with information (such as operating system and host criticality) about the systems affected and, ideally, prioritized based on this additional information. For example, it's much more important to know about an attack on a system vulnerable to a particular exploit than it is to know about the same attack on a patched system that is otherwise not vulnerable.

Ideally, this big picture context is something that can be amalgamated from other security systems (such as patch management and vulnerability scanning systems) and external databases (such as asset management systems and user directories); in a more sophisticated environment, this should be augmented with information that the IPS can supply.

To bring this necessary context to their IPS sensors, some enterprises have turned to the SIM (Security Information Management, sometimes also called SEM, for Security Event Management) market. SIM vendors have made a good business of providing the context and analysis tools that many IPS products lack. However, for best results, the analysis console needs to be very tightly tied to the IPS product itself. For example, one of the common results of IPS analysis is a change in IPS policy. The enterprise manager wants to be able to effect policy change with a few clicks in the same interface, something you can't do in a third-party SIM tool.

4) IPS must also offer IDS-like features both in management and in the sensors

IPS and IDS products have very different requirements for performance, analysis, and accuracy. IPSes must be high speed and low latency, are primarily designed to block attacks, and should be tuned for the lowest possible level of false positives. In contrast, IDSes do not have to be engineered for peak traffic levels, are primarily designed to discover and report

on attacks, and may have some false alarms and false alerts without affecting their usability or the enterprise network.

Nevertheless, in an enterprise, every IPS is expected to also have several IDS features. In fact, the best IPS products incorporate both an IDS and an IPS in the same package so that the network manager can both have the blocking power and performance of an IPS along side more intense alerting and deep analysis features traditionally built in an IDS.

Simply put, any enterprise IPS should be an IDS as well, able to both detect and block threats in the sensor, and with the analysis and forensics management toolkits commonly associated with an IDS.

The attention paid to IPS alerts is often lower than to IDS alerts, but this does not mean that an IPS can get by with a lower quality of analysis tools. When something goes wrong on a network, the alerts out of an IPS may be the only clues available to what happened, who is to blame, and how to most expeditiously repair the damage. This makes the analysis tools for an IPS, even if seldom used, just as critical as the tools for an IDS.

Although blocking and detection functions are very different, there are commonalities between IPS and IDS that make co-location of the function very appropriate. For example, the detection engine used in an IPS is often the same engine that is used within a standalone IDS. Because of the natural affinity of these two functions, enterprises should expect that they can get both IPS and IDS functionality out of the same box.

There are other reasons to get both functions out of a single box as well. With highly switched networks, it is becoming more and more difficult to place IDS sensors within the enterprise, since monitoring ports may be difficult to provision or may already be occupied by other security appliances. It's also undesirable to keep propagating security devices around the network. By having IDS and IPS in a single device, you cut down the total number. While performance concerns may preclude this co-location of function in some environments, it is a natural conjunction in others.

The best enterprise-focused IPS products should allow for "virtual sensors:" the ability to break up the physical data stream into pieces, separated by LAN segment, source IP or destination IP, each of which is evaluated according to a different security policy. This allows for different sets of IPS and IDS features to be brought into play different parts of the network, all contained within the same physical device.

5) IPS vendors must offer centralized management over multiple devices

An enterprise IPS/IDS deployment never consists of a single standalone device; it could include tens or even hundreds of sensors located across a large LAN or even larger WAN. The burden of managing these security devices must be minimized and streamlined so that security staff is able to concentrate on what matters—the information that comes out of these devices—and not what doesn't—the mundane task of keeping them up-to-date and operating properly.

Some of the key requirements for a good centralized management system are already summarized in the "big picture context" requirement described above. A few others are worth mentioning because of their relative absence in the marketplace. These include overall system management, information aggregation, and the ability to subdivide the task of security analysis across multiple teams.

The management system for an enterprise IPS must be proactive rather than passive; it's not simply a matter of collecting alerts and providing a pretty dashboard. The management system has to look for problems (such as impending performance bottlenecks) and keep IPS signatures updated and fresh on all managed devices. While we're accustomed to having this class of centralized management for systems such as Windows desktops—indeed, we probably couldn't live without it nowadays—there are many IPS and IDS deployments without a strong central management system. Without efficient centralized management, these deployments are the ones least likely to succeed in either blocking legitimate attacks or providing the evidence needed to detect security problems on a network.

The management system also has to integrate alerts and information from multiple devices simultaneously, summarizing duplicate information, and presenting the security posture of the network as a holistic entity, rather than a bunch of small bits broken up into arbitrary buckets. Enterprise security analysts can't properly interpret and manage the multitude of alerts that will flow in from network of IPS/IDS devices without the assistance that a good management system offers. The centralized management system has a vital place in providing sufficient, aggregated and organized information to support the analyst in their job.

A final requirement of central management is the ability to separate out management functions. In an enterprise, management of an IPS often needs to be divided both horizontally and vertically. Horizontal separation lets different

functions be controlled by different teams, such as separating operations and monitoring tasks from policy updates. Vertical separation keeps the sensitive security information within the management system compartmentalized, so that different teams see only the information that is relevant to their jobs. This slicing of the management system lets security analysis and threat management be distributed, safely, across different parts of the organization.

6) Offer the agility to quickly translate security information into security policy

Any IPS or IDS operates using a policy: what to block, what to allow, what to alert on, and what to limit to a certain rate. The process of properly creating that policy is difficult enough that this task, alone, is responsible for many IPS and IDS failures. A security device that cannot be controlled by policy does not meet the needs of any network, large or small.

A critical requirement in enterprise IPS is the ability to easily and quickly manage and control policies on the deployed sensors. The centralized management system must be able to slice and dice the information coming from the sensors so that the network manager can understand what it is saying quickly and easily. And, just as importantly, the understanding that the security analyst gets from the console has to be just as quickly and easily fed back into policy.

Enterprises that have selected external SIM devices to act as their IPS "super-consoles" should be acutely aware of this problem. While some of these SIMs offer enticing capabilities to correlate and aggregate information collected out of IPS devices, they have a critical flaw: they can't be used to manage the policy on those devices.

In an enterprise IPS deployment, the security analyst needs to be able to point at an event and, within a span of two or three seconds, tune the policy that caused that event to fire. That may mean suppressing event alerts for a particular system, adding an IP address or subnet to a "trusted system" list, disabling a policy, changing the policy action (such as from "drop" to "alert"), or even modifying rule parameters or values. Without a quick feedback loop that empowers the analyst to tune quickly and efficiently, the fallback tuning mechanism is going to be turning the IPS off entirely (or changing it from an IPS to an IDS, thereby losing much of the value of an IPS).

This quick feedback loop between analysis and policy has other implications for enterprise-class IPS. The need for "virtual sensor" (discussed above), gives the analyst the ability to deploy sensors and policy finely tuned to the systems it is protecting. Policy controls themselves must also be fine-grained, allowing easy management and maintenance of exception and suppression lists, without having to turn off protection behaviors just to get the IPS to behave properly. And the management system itself must be designed to let the security analyst quickly determine which controls and adjustments to policy are necessary, and just as quickly apply those adjustments and re-deploy policy.

7) Simple firewalling and other network security tools to help in branch office deployments.

In many cases, IPS and IDS sensors are obviously standalone devices that should be separate from firewalls and placed at very different parts of the network. However, there are also equally valid deployment scenarios where IPS and firewall functions should be co-located at the same point. An IPS is not a firewall, but the policy within an IPS is very directly derived from the expected traffic at any point on the network—an expectation that is enforced by a firewall.

Enterprise security architects want to reduce the number of devices and policies they have to manage, which means that having some simple firewalling capabilities in an IPS (just as some firewalls have taken on simple IPS capabilities) is desirable. Having a combination of capabilities can make the difference between an awkward and difficult-to-support configuration and one that is both efficient and elegant.

This doesn't mean that the ideal IPS is a UTM firewall: the IPS capabilities in UTM firewalls are generally very restricted and aimed more at the SMB environment than an enterprise network. However, in several environments, especially branch offices (where device count, cost, and heat/power/noise are significant evaluation criteria), having a "small firewall" inside of the IPS can make the difference between a successful deployment and an unsupportable mess.

Wrapping Up

This white paper has focused on specific technical requirements that an enterprise-class IPS must meet, including:

1. Availability of a broad hardware range with offerings ranging from branch-office T1-class devices to enterprise multi-gigabit systems;
2. A detection engine incorporating multiple technologies, including signature matching and protocol anomaly and behavior anomaly detection;
3. An ability to provide a "big picture" context of information about the network and its devices to qualify security information as well as speed analysis;

4. IDS-like features to detect less critical threats, perform behavioral analysis and facilitate security forensics toolkits;
5. Centralized management and alerting capabilities across multiple devices;
6. An analysis toolkit that efficiently drives policy refinement; and,
7. Network security tools, especially firewalling, to simplify deployment

Other "soft" requirements are, of course, equally relevant.

- An IPS must be designed to merge smoothly into the operations requirements of the enterprise, making both the sensors and management system more than just simplistic black boxes. Features such as SNMP monitoring and alerting, off-box configuration backups, and log file archiving and rollover fit into enterprise operations procedures cleanly.
- Without constant updates, an IPS turns into a doorstop. This means that evaluating financial stability and long-term commitment to the IPS market of the vendor is a key part of the purchase decision.
- Because technical support is important in helping the analyst understand and control the IPS, a good support infrastructure with 24x7 coverage by security professionals (not an answering service) is another requirement.

Enterprises have shied away from IPS products because, honestly, the vendors couldn't meet the needs of the enterprise network managers and security analysts. The time has come to re-evaluate the place of IPS in enterprise networks. Using the requirements in this white paper, you can identify IPS products that are finally ready for enterprise deployment.