Six Integral Steps to Selecting the Right IPS for Your Network

Joel Snyder

Six Steps for Selecting the Right Network Intrusion Prevention System to Suit Your Network

About the Author: Joel Snyder, PhD, is a senior partner with Opus One, a consulting firm in Tucson, Arizona. He's helped over 200 private and public organizations with their networking, email, and security problems. He spends most of his time on the road helping people build larger, faster, safer, and more reliable networks.

Joel has been working with networks and information security since 1981, beginning with X.25 and public key cryptography and he's covered most topics in networking, electronic mail, and information security since then. He has been a member of the ISO and ITU committees which write network standards, has authored several books and hundreds of articles. As an author and speaker, he's received numerous awards, recognizing his work to improve enterprise IT.

Executive Summary: Network Intrusion Prevention Systems (IPS) can be extremely effective pieces of your overall network security strategy. However, the IPS marketplace is filled with products that all do very different things and are suitable for very different environments. Therefore, buyers beware, because simply throwing any IPS into the network without careful consideration can be a costly error, both in terms of capital outlay and operational provisions.

The critical question to answer is: "Why are you buying an IPS?" (Step 1) Answering this question will help define both what you want in an IPS and help you weigh what you can expect to get from these products as you evaluate them for use in your network.

With the answer to this underlying question in hand, you'll be well positioned to closely examine four aspects of IPS products that distinguish them from each other:

- Security parameters and coverage (Step 2)
- Performance (Step 3)
- Form factor (Step 4)
- Management (Step 5

Considering how each product delivers on these four characteristics will allow you to quickly and efficiently create a short list of products that you will need to evaluate and test in your own network (Step 6) as the final –and essential - part of assuring that you are achieving the goals that will justify the costs associated with deploying an IPS in your network.

It's a dire fact that while every enterprise has a firewall, most still suffer from network security problems. IT professionals are acutely aware of the need for additional protective technologies, and network equipment vendors are anxious to fill in the gap. Intrusion Prevention Systems have been promoted as cost-effective ways to block malicious traffic, to detect and contain worm and virus threats, to serve as a network monitoring point, to assist in compliance requirements, and to act as a network sanitizing agent.

While all of these capabilities may fall within the purview of an expensive, high-end IPS product, not every IPS deployment will require all of these features nor will every business be able to accommodate the operational price necessary to maintain and manage one of these high-end systems.

For these reasons, the IPS market is overflowing with products that are suitable for a wide array of environments as they offer a wide spectrum of features.

Establishing which IPS product is right for your network is crucial to any buying decision, because putting the wrong IPS into your network can be a costly error, both in terms of capital and operational expenditures.

In the IPS world, it is especially easy to fall into the trap of buying what a particularly savvy vendor wants to sell you, rather than what you actually need. In this white paper, we'll provide a six-step strategy for deciding what IPS is right for your network that begins by answering the question "Why am I buying an IPS?" and ends with a plan for testing an IPS in your own network.

STEP 1:
Answering the question "Why am I buying an IPS?"

The most critical step to making good decisions about which security products to implement throughout your network is to first know what it is that you want to accomplish. Before looking at products, before talking to vendors, and certainly before deciding whether you even need more security, you need to answer one simple question: "Why am I buying an IPS?"

There are many good reasons to add an IPS into a network. You could be looking for extra protection at the perimeter, something that faces towards the Internet and employs signature-based technology to trap some of the bad things that manage to make their way through the firewall. Or, you could be more focused on mitigation of denial-of-service attacks, and looking for products that employ rate-limiting security parameters to protect against these kinds of threats. With a new, onerous, load of regulation in many organizations and industries, you could be looking for tools to help in your compliance efforts. Or, perhaps you might be looking for a product that provides IDS-like alerting and forensics to help you get a better handle on just what kinds of threats are trying and have been successful at hitting your network.

You could be hoping to build more security into the core of the network, perhaps protecting a specific set of servers inside the network or even by wrapping an IPS around the entire network core. You could be worried about incoming threat—or just as worried about detecting and blocking infected systems on your own network from attacking the rest of the world.

Note that this isn't a comprehensive list; it should simply be used as a start to some conversations about the possible reasons driving your decision to add more protective technologies to your network.

Each of these reasons noted above can be equally valid in the right environment, but until you know which ones apply to you, you won't be able to select the proper IPS strategy or the product that will best realize it. Every IPS has a different set of design goals and features targeted to address a limited set of the questions posed here.

It would be easier for all involved if you could simply reduce this list of implementation reasons and goals into a feature checklist, something you could throw into an RFP and subsequently pick the vendor with all of the right boxes checked. But, unfortunately, that's impossible, not so much because the appropriate features are not in place, but because of the disparate philosophies that go into the products' design.

For example, it's easy to put "forensics" onto your checklist as a featureassuming that is something you care about. Unfortunately, listing "forensics" won't get you any closer to finding the right product; it will only help you to eliminate some products that don't have any forensics capabilities at all.

The more appropriate question is why do you want forensics? Are you really looking to comply with the classic definition for forensics in which you need to collect data that could be used in a courtroom to help prosecute an attacker? Or are you simply looking for data collected and stored over a period of time that will ultimately help you to understand how an attack actually happened? Will you need to tap into the forensics ability of the IPS daily or just once a month? If you expect to run daily forensics, the performance and design of the forensics interface is a huge issue, while they may not be as important if you only need to review on a monthly basis knowing why you want forensics will help you to understand what products will work best for you.

This issue of truly understanding why you're adding intrusion prevention and what you're looking for in IPS is so critical that it's difficult to under-emphasize its importance. The IPS market is crowded on many levels. There are products ranging from high-performance standalone appliances to ones shipped as add-ins to existing firewalls. After studying this product space for several years, it has become clear that while there are often common denominators between some products –for example, quite a few of the newer IPS products use Snort as their underlying detection engine– that help segment the market into broad, overlapping categories, the underlying design goals and capabilities still vary widely.

Table 1 comprises a list of many of the reasons why corporations we've worked with in the past three years have implemented an IPS in their networks and the noted tradeoffs expected with each choice. Use this table to help guide you to make your own IPS needs statement. No single IPS device is designed to operate in every environment and solve all problems, which means that you will have to make choices and weigh your own reasons to balance these tradeoffs.

TABLE 1

TITLE: Contrasting Reasons for Adding Intrusion Prevention Systems

Spectrum of Reasons for Implementing IPS		Design Characteristics of an Appropriate IPS
From:	To:	
You are focused on perimeter security	You want to protect the core of your network	The closer an IPS is to the core of your network, the more important issues such as performance, high availability, and control of overflow become. IPS functions pushed out towards an Internet boundary don't necessarily operate under the same performance constraints, and may be designed to handle failure cases (such as too much traffic or too high latency) differently.
You want to protect your servers	You want to protect end users (clients) on your network	When protecting servers, an IPS can be tightly tuned to inspect particular incoming services and particular applications. To protect client desktops, the IPS must both handle incoming and, more importantly, outgoing traffic with twin goals: prevent incoming infection and attack by blocking packets, but also detect a compromised system by its outbound attacks.
You are looking for signature-focused IPS protection	You are looking for rate-focused IPS protection	While most IPSes have both signature-based and rate-based technologies in place, one or the other is generally the "sweet spot" for the product. For example, when your main concerns are denial-of-service attacks, a product architecture focused on rate-based IPS is needed. If you are more focused on break-ins through system vulnerabilities and reconnaissance, signature-based IPS is more appropriate.
You are most concerned about specific attacks, such as hacker break-ins or viruses.	You are most interested in detecting anomalous behavior, such as a normally unused server	Although these two capabilities are by no means exclusive, most products specialize in one or the other. Simple anomalies, such as protocol errors, are common across the board (even in rate-based products), but more sophisticated detection scenarios, such as behavior anomalies, call for a different architecture.
You want to be able to detect attacks and have some forensics evidence on how it happened.	You want the IPS to operate on its own, but you are not interested in using it as a security console or as a primary tool in investigations	While an IPS can both detect and prevent attacks, adding a full forensics capability of any sort dramatically changes product architecture, increases costs, and impacts performance.
You want IPS in place for primary protection against attackers and break-in attempts	You want IPS as an additional layer in a Defense in Depth strategy	IPS products positioned as a primary protective layer, typically behind a firewall, may have other features such as "shunning" of known attackers. These bring additional security, but at considerable risk such as self-inflicted denial-of-service . When an IPS is part of a layered defense strategy, features such as shunning are often unnecessary.

To understand why you're looking for an IPS, write an IPS needs statement, a single paragraph that begins with this phrase: "What we're trying to accomplish is …" With this in place, you'll be in a much more informed position to correctly evaluate IPS products for your environment. Only after you understand **why** you want to add an IPS to your network can you ask yourself about security and coverage, performance, management, and form factor—the other four main criteria for successfully selecting an IPS strategy for your network.

STEP 2:
Determine the Level of Security and Coverage you require.

The term "Network IPS" doesn't inherently imply any one way of preventing intrusions. In fact, different products use radically different technologies to help add security to networks—because "security" means radically different things to different people. Your voyage down the IPS path will be smoother if you take a moment to examine your own definition of the security you hope to gain from an IPS and state it explicitly. Are you talking about integrity and availability of your network? Regulatory compliance? Application security? Leak protection? Each of these could be a valid component, but getting your team on common ground is a good step at this point. At the same time, you should consider coverage—what kinds of protocols and applications will the IPS be responsible for inspecting and understanding, and at what depth.

While the vendors' marketing departments make great efforts to distinguish the underlying technologies, there are fundamentally three approaches in current IPS products: signature-based (including protocol anomaly) IPS, rate-based IPS, and behavioral IPS.[2] While the leading products may include some pieces from all three approaches, each product has a fundamental direction it follows—either signature-centric, rate-centric, or behavior-centric—with the other two approaches being secondary and tertiary.

The important part of this step is to decide which of these three approaches is most important to you overall and most appropriate for your application (Refer to your "Why" statement here for continuity.) Each approach to Intrusion Prevention gives a different kind of security and a variant level of protection, and sits in a different spot on your network.

The dominant form of IPS in the marketplace is **signature-based** IPS. These products are readily available and range from remotely managed service-based devices to standalone high-performance IPS to embedded IPS technology in firewalls. Signature-based IPS products do not rely entirely on signatures to detect malicious or improper behavior. Many also include other detection technologies. For example, a detection technology good at catching "zero day" attacks is protocol anomaly detection, which looks for application or TCP/IP behaviors that are either non-standard or far from the normal behaviors (such as an SMTP "recipient" address with 500 characters in it, or TCP packets with malformed options in them). Most signature-based products will include some protocol anomaly measures in their repertoire as a means of thwarting zero day attempts.

Signature-based IPS technology is critical to catching and blocking common exploits, but it's also important to understand that signature-based IPS has significant limitations. A signature-based IPS is only as good as its signatures, and writing signatures is a difficult art, made still more difficult to evaluate since very few vendors actually offer open signatures which can be inspected. Although a mantra of signature-writers is to "block the vulnerability, not the exploit," the reality is that many IPS signatures are only good at catching well-described exploits and do not necessarily protect against the underlying vulnerability. Because most systems see many different data streams as equivalent, long considered a desirable attribute of a well-designed and interoperable Internet application[3], many IPS signatures have an Achilles heel in their inability to identify every possible permutation of an attack that will exploit a vulnerability.

Even with all these technologies brought to bear, most signature-based IPSes are best at detecting use of common exploits (for example, by attackers simply trying tools they've downloaded from the Internet) and not as capable in blocking a true, targeted attack. If your main worry is attacks that might exploit unpatched and unprotected systems, signature-based IPS will block the script kiddies' attempts to compromise your systems, but not someone who has insider information and is intentionally trying to evade the IPS.

[1] The concept of "self-inflicted denial-of-service" is a common one used to argue against automatic blocking (commonly called traffic shunning). The premise is that unless blocking is done very, very carefully, it's easy for a system to block access to vital resources. The common example is the Internet root DNS servers. If traffic to these were blocked, outbound access (and some internal traffic) will be very quickly restricted, a self-inflicted denial of service caused by not being very careful in automatic blocking.

[2] These last two (rate-based and behavioral) are often thrown into a single category of "anomaly-based" IPS or called NBAD for Network-based Anomaly Detection. In fact, they have very different characteristics and should be considered as separate technologies.

[3] "In general, an implementation should be conservative in its sending behavior, and liberal in its receiving behavior," from the IP standard, has become such a mantra that this Robustness Principle appears in nearly 50 RFCs.

Rate-based IPS works by closely watching the rate at which connections come into high-performance application servers, most typically Web servers. The primary goal of rate-based IPS is to mitigate and protect against denial-of-service attacks (whether intentional, or unintentional, as misbehaving software might be a likely root cause). Rate-based IPSes are definitely in-line devices, because they take an active part in monitoring, controlling, and filtering connections. Rate-based IPS products can both detect simple overloads (such as too many connections over a short period of time, typical of a Botnet originated DoS attack or a particularly popular slash-dot posting) as well as attacks based on half-open connections (such as those that try to fill application server process tables or firewall state tables with incompletely established connections).

The best rate-based IPS will actually step in and shield servers from bad connections during periods of stress by proxying connections to be sure that there is someone 'alive' on the other end. More sophisticated rate-based IPS, appropriate for huge application server farms, offer a myriad of fine-tuned controls, but the basics of rate-based IPS can be built into any in-line IPS device or firewall. These technologies scale down very well and can easily protect small and medium-sized businesses with Internet-facing servers from many types of denial-of-service attacks.

Since rate-based IPS is best aimed at the perimeter of the network, embedding the technology into firewalls is the most appropriate strategy for all but the largest of data centers. Again, the IPS needs statement created in Step 1 of this process will help to determine whether rate-based IPS is your primary requirement, or whether it is an adjunct to other intrusion defenses.

Behavioral IPS tracks the flows and traffic patterns of a network.[4] When these change, the IPS alerts the security manager and, in extreme cases, blocks or throttles traffic. Behavioral IPS is poor at detecting or blocking specific incoming attacks because most attacks, based on a specific data stream embedded in a normal protocol transaction, are not actually changes in behavior. However, these systems are very good at identifying systems that have become infected and are now attacking other systems and users, or which have become bases of operation for hackers.

Behavioral IPS offers an interesting view to network managers, especially in large, complex networks where the actual flows are not fully understood as a general rule. For that reason alone, many behavioral IPS systems have become valuable tools. However, behavioral IPS is barely comparable in its ability to actually prevent intrusions to rate-based and signature-based IPS, and solves very different problems. You should already know whether behavioral IPS is appropriate for your network based on your IPS needs statement outlined in Step 1 of this process.

It's impossible to say which type of IPS offers the "best" security, because each of the detection technologies has different characteristics and helps in different ways. What is important is matching the type of security offered by the IPS with your requirements as outlined in your IPS needs statement.

Once you've decided the type of IPS offers you the best capabilities for your requirements, a second decision comes into play: what coverage do you want? For signature-based IPS, the quality (much more than the quantity) of the signatures and the underlying detection engine determines whether the product will meet your needs and is right for you. Even in signature-less systems, such as rate-based or behavioral IPS, coverage of applications and protocols will vary from product to product.

Some IPS products have evolved from Intrusion Detection engines, and are sometimes glibly referred to as "IDS with the IPS bit turned on". In evaluating IPS products, you may run into ones with literally thousands of signatures, a sign of an IDS that has been repurposed as an IPS. An IPS's job is not to detect every possible attack or reconnaissance attempt; an IPS only has the job of preventing intrusions. That job takes considerable smarts. The signature count of a well-designed IPS engine will number in the hundreds, not thousands. This doesn't mean that an IPS cannot have additional signatures and also serve as an IDS, as long as the IPS function is primary and the IDS function is secondary.

However, an IDS which has simply been put into inline mode with its thousands of signatures will offer a lower degree of security against serious attackers than an IPS designed from the ground up as a prevention, rather than as a detection, system. For example, because the IDS signatures are designed to detect attacks, even attacks which might never succeed, they will have a higher false positive rate than IPS signatures that are designed to identify and block working attacks. At the same time, the burden of having thousands of signatures to inspect traffic across many ports and in many directions will cause performance issues (such as high latency or last packets) that might be acceptable in an IDS, but would never be allowed in an IPS.

[4] Behavioral IPS is often considered IDS, intrusion detection systems, because they rarely even attempt to prevent intrusions. However, as the market hype for IPS has eclipsed IDS, many vendors of behavioral IDS products have either modified them to have some blocking capabilities, or simply rebadged them as intrusion prevention, rather than intrusion detection, systems.

Merely taking an IDS device from a monitoring port and moving it directly into an in-line IPS position without significantly re-engineering the system is a recipe for failure.[5] Neither the IDS detection engine nor the IDS signature set is very appropriate for an IPS deployment for two reasons. First, there are fundamental differences in design philosophy between IDS detection engines and IPS detection engines. Secondly, simply repurposing IDS signatures into an IPS will create a fairly inaccurate IPS. For example, an IDS should behave differently when it alerts depending on the susceptibility of a destination system. That is, an IDS might handle a Code Red attack on a Unix server (which is not vulnerable) differently from a Code Red attack on a Windows IIS server (which is vulnerable). In short, it should handle a successful attack very differently from an unsuccessful attack. However, the IPS doesn't need to make any of those judgment calls. It simply has to block Code Red attacks altogether, a significantly easier job.

On the other hand, an IPS doesn't always have the simpler job. Fundamentally, the penalty for a false positive in an IDS scenario is low, which means that there is an incentive to err on the side of producing false positives. Conversely, the penalty for a false positive detection in an IPS deployment is huge, which means that the IPS designer must have a significantly different mindset.

For example, when the team at Oulu University in Finland discovered widespread security vulnerabilities in common SNMP implementations, IDS vendors created signatures that quickly matched the exploits used by the Oulu PROTOS project. However, detecting an exploit is not the same as detecting vulnerability. Our testing has shown that these SNMP signatures both have a high false positive rate in that they detect valid uses of the protocol and mark them as exploit attempts, as well as a high false negative rate in that they fail to detect very simple and slight variations that a moderately skilled attacker might use to evade an IDS.

In this situation, an IPS has the more difficult job because it must not block valid SNMP calls, yet it should detect attempts to exploit these vulnerabilities. Here, the IPS must combine the use of protocol analysis with signature detection. Protocol analysis is also invaluable in providing "zero day" coverage to help catch exploits against previously undiscovered vulnerabilities.

Finally, when considering security and coverage, you should look at the potential IPS actions. A simple IPS can only drop offending packets, but more sophisticated actions are also available with more advanced products. These actions include resetting connections (in one or both directions), sending alerts, capturing packets, blocking future traffic, and even changing configuration of other network devices, such as firewalls. While the more advanced actions, such as blocking future traffic (sometimes called "shunning" after Cisco's terminology) or changing firewall access control lists, may seem attractive, our experience is that these cause more problems than they solve. Thus, making these capabilities a requirement in picking your IPS could be overkill, and could force you to disregard perfectly capable devices.

STEP 3:
Determine Your Performance Requirements.

IPS performance is something you can't afford to get wrong. Unfortunately, performance of IPS devices is difficult to test, and the results are almost as hard to describe. As products like IPSes move further up the network stack, their performance becomes highly data-dependent. This is different from what we're used to witnessing in the world of switches and routers, where performance is easy enough to describe. Even for firewalls (at least firewalls without UTM features) performance is easy to measure because metrics such as connection rate, maximum simultaneous connection count, and goodput (often called throughput) are commonly understood and universally accepted.

IPS devices are much harder to characterize. The greatest differentiator in performance is not the IPS itself, but how it is configured. For many signature-based IPS products, the performance of the product varies hugely based on the number of signatures and protocol decoders enabled for detection. For example, an IPS may have hundreds of signatures covering HTTP. If half of those signatures are disabled (perhaps because they are IIS signatures and Apache is being used), then the performance of the IPS on HTTP traffic can be quite different. Similarly, many IPS vendors classify signatures by severity. If only "high priority" signatures are enabled, the IPS will pass traffic more quickly than if all priority types are enabled.

Your traffic may also cause variations in performance. For example, an IPS may be able to pass clean HTTP traffic at 2Gbps rates---unless the traffic is in Japanese, at which point the rate can drop to 1.75Gbps. Why? Asian languages use multi-byte characters, and the HTTP processors inside the IPS have to do much more work with multi-byte HTTP. More commonly, an IPS will have dramatic performance

[5] We assume that all IPS devices are in-line devices, since that's what's required to actually prevent intrusions. Devices that claim to have IPS capabilities yet are not in-line to network flows may have desirable security characteristics, but those are focused on detection and reporting rather than prevention. Any signature-based or rate-based IPS that does not require in-line operation for prevention should be regarded with extreme skepticism.

differences based on the protocol used to pass the traffic. For example, moving files around a network with Windows file sharing might not slow down the IPS very much because there aren't many IPS signatures for Windows file traffic. If you moved the exact same files using a protocol that has more signatures and requires work to decode and normalize traffic, such as SMTP, you would see very different performance characteristics.

Additionally, IPSes will also behave differently depending on the mix of attack traffic and benign traffic.

In our testing, we found that attack traffic has a disproportionate impact on IPS performance compared to "clean" traffic. Because an attack is considered an exception, has to be logged, generates an alert, and generally requires much more processing than non-attack traffic, the ability of an IPS to pass traffic as the attack rate goes up varies dramatically with small amounts of attack traffic.

If you intend to put an IPS out near the perimeter of your network, you will see more attacks—and thus greater variation in system performance. The worst performance case would be to put an IPS outside the network firewall, fully exposed to the Internet. This has the advantage of providing the curious security staffer hours of amusement and gigabytes of interesting data. It also has the downside of slower and generally unpredictable performance because of the variability in type and volume of Internet-sourced attacks.

As an IPS moves closer to the core of the network, the ratio of attack traffic to normal traffic will change so that observed performance become much more consistent. While an IPS protecting internal systems does have to handle a very high transaction rate, much higher than one simply at the network perimeter, it will also see a smaller amount of attack traffic.

A critical step before adding any IPS to your network is validating the vendor's performance claims by testing in your own network, using live traffic, and using your selected signature set. In published benchmarks, traffic may have been hand selected to be "low impact" on the IPS, and a minimal set of signatures and decodes turned on. This may make good marketing literature, but it represents a dangerous way to specify the performance of IPS devices.

To determine the real performance in your network, make sure that the protocols you use and the signatures you care about are all enabled. This may require some amount of tuning on your part, but it's better to discover performance limitations before committing to a full IPS deployment. More details on IPS test methodology are in Step 6 of this white paper.

A second aspect to IPS performance lies within the management system. If your goals for implementing an IPS call for forensics functionality or alerting and reporting, testing the performance of the management system should be part of your evaluation process. Our IPS testing has shown that many IPS management systems will slow to unacceptable performance levels when more than a small number of events are arriving in short period of time, or when a significant number of events have accumulated in the management system database. With IPS devices being pushed as part of regulatory compliance, where years of record keeping are generally required, performance of the management system with millions or tens of millions of events requires some validation.

STEP 4:
Determine Your Form Factor Requirements.

IPS is not a product; IPS is a function and a technology. You can package that technology in many ways, and place that function within many kinds of devices—including standalone IPS appliances, inside of firewalls and switches, and in other types of security appliances, such as SSL VPNs. When you consider IPS for your network, your choice of form factor (appliance or integrated function), and where you will place the IPS function in your network will dramatically affect the products you should consider.

Unfortunately, it's not easy to divide IPS functionality strictly along the lines of form factor. While standalone IPS appliances offer a very high level of IPS functionality, it doesn't mean that an IPS integrated into a firewall or switch always has a lower level of security, coverage, and performance. At the core of a network, standalone IPS products will probably be the most appropriate to meet performance requirements and keep topologies simple. But at the edge, IPS integrated into a firewall may be the best form factor choice. Embedding an IPS in another device, such as a firewall, brings its own complication, because now you must evaluate the quality of each component. For example, an IPS with excellent capabilities integrated with a poor quality firewall is a poor compromise. In fact, the interest of high-end firewall manufacturers in bringing IPS technology and functionality to their customers means that a few firewalls have integrated IPS functions (usually delivered by adding hardware into a chassis-based system) that offer the same functionality as standalone devices. Even with mid-range firewalls, vendors have brought in sophisticated IPS functionality, usually focusing on protocol anomaly detection and a small set of signatures that may be sufficient for your requirements.

On the other hand, some firewalls have an "IPS function" which was placed into the device simply to satisfy a checklist requirement as part of a Unified Threat Management (UTM) offering. In almost every case, these IPS features are based on some version of the Snort IDS engine, with the Snort signature set either included in full or trimmed up by the security vendor. Although Snort does a poor job as an IPS--it was designed as an IDS and its detection technology and operation is not optimized for intrusion prevention--this isn't the main reason why these embedded IPS functions in UTM firewalls should be avoided.[6]

The real problem with embedded Snort-based IPS in UTM devices lays in system management. Because Snort currently has over 6,000 detection rules (with an additional set of "Bleeding Snort" rules that are even more important in detecting recent attacks), the burden of deciding what traffic should be subject to the IPS, which rules should apply, and what the action should be, is an enormous prospect. More importantly, when the inevitable alerts -- and especially false positives -- occur, a typical Web-based interface isn't going to be up to the task of helping the security professionals figure out which signature was triggered and which needs to be disabled for which traffic. The result of this complexity is that the security professionals are never able to effectively configure the IPS to add security, while keeping the false positive rate at an acceptable level. The vast majority of UTM firewalls with Snort-based IPS functionality have the IPS disabled, as is appropriate

Fortunately, not all firewall vendors have chosen to take the easy route and put in a poor IPS just to meet a specification. Once you've discarded the bad UTM firewalls, this still means that you have to make a decision: what is the form factor most appropriate for my requirement as outlined in the IPS Needs Statement detailed in Step 1? The three most common options are a basic IPS in a firewall, a full IPS co-located in a firewall chassis, or fully freestanding IPS.

Basic IPS in a firewall, typically focusing on behavior and protocol anomalies, is an excellent choice if you have a good patch and security management policy in place on all internal servers, specifically those accessible from the Internet. In that case, the additional layer that an IPS offers on top of existing firewalls and well-maintained systems is some protection from day-zero attacks as well as denial-of-service attacks. Although no vendor can promise true "day zero" protection, basic behavior and protocol anomaly, as well as simple rate-based controls, add a huge amount of value in their capabilities to block common attack methods and protect servers against traffic overloads on top of a normal firewall.

Full IPS in a firewall is the best strategy if your main concern is Internet-sourced attacks and, to some extent, identifying internal systems that have become infected or compromised. The benefits to network topology and operations costs of putting the IPS within the choke points of the network are great. They reduce the complexity of the network over the alternative of a standalone IPS sitting next to a firewall, which thereby increases reliability. At the same time, having a firewall and IPS co-located in the same system offers opportunities for management that standalone boxes cannot easily support. For example, the firewall could only send a subset of traffic through the IPS, speeding performance and eliminating the possibility of false positives in critical environments. Since the firewall rules and IPS rules are synchronized within the same system, the IPS can "know more" about the traffic and make better prevention decisions.

Standalone IPS products are most appropriate in two environments. Most obvious is when the goal of the IPS is to protect a set of systems from both external and internal threats. By pushing the IPS closer to the systems being protected (rather than the Internet), the IPS protects against all attackers. The second environment where standalone IPS is appropriate is one where IPS and security auditing are organizationally divorced from firewall configuration. For example, in some organizations faced with regulatory compliance issues, IPS and IDS tools are managed by a separate audit group, one that is organizationally separate from the security operations team.

STEP 5
Determine your Management Requirements.

Management of IPS is a huge issue in product selection, and matching your requirements for management, monitoring, and forensics capabilities with the product you choose is as important as any other selection criteria. IPS products vary in their management philosophy from "virtually no continuing management" to "very high management requirement" styles. These management styles reflect not only the philosophy of the product design team, but the configuration needs that any design implies. A mismatch between IPS management requirements and the product you select can lead to catastrophic failure of your IPS deployment. The worst thing you can possibly do is select a "high management" product and put it into a "no management" environment.

Many IPS management systems are unlike any other application or management system in the network. This difference, and the accompanying complexity, is an important

[6] I am trying here to avoid the semi-religious argument about whether Snort is a good IPS. I don't believe it is, but whether you agree or not, this isn't the reason that Snort makes a bad IPS when integrated into a UTM firewall.

factor, especially if you don't have the luxury of a dedicated IPS/IDS team. As you determine management requirements, keep in mind who will be responsible for day-to-day management of the IPS, what their level of expertise is, what more they can be expected to learn, and how many hours a day you've budgeted for IPS management.

Some of the other factors that will affect your management requirements include forensics needs, event alerting and lifecycle needs, and performance needs.

Forensic capabilities come about because many IPS products also have IDS capabilities. Although simply turning an IDS into an IPS doesn't give you a good IPS, having an IPS with a lot of IDS features in place can bring a lot of value to a security analyst. This type of feature set—intensive logging, inclusion of IDS signatures, and packet capture are three key indicators here—is an early decision in your IPS deployment plan. As a security analyst, I believe that IPS products with this type of capability are a great addition to any network, contributing to network understanding because it gives you the ability to look at security problems after-the-fact. In some cases, an IPS with IDS features can even replace a standalone IDS.

However, it's important not to look for IDS and forensics capabilities if you don't intend to use them. The cost of maintaining a high-speed management database for IDS is high, as is the amount of hardware and maintenance required to keep such a database running. Paying for a high-end management server that can store a year's worth of alerts and their forensic information is only OK if you actually want to use it. Some IPS products are flexible enough to support either mode of operation: with packet captures and forensics, or without. If you're uncertain what your IDS and forensics requirements are at this stage, you should consider specifying a device that can operate just as easily with packet captures on or off.

Network visibility is a valuable side benefit from many IPS products. Because they see so much traffic, they can provide both network and security managers' insight into what is happening on the network. IPS management systems that present this information graphically offer great benefits and can highlight problems at a glance---which makes basic activity analysis easier.

Event alerting and its correlating event management capabilities are a second set of management features that can differentiate IPS products. For some IPS devices, the only goal of alerting is to provide a brief track-back to help eliminate false positives. These products may store a few days of alerts and have limited capability to search and manage these alerts. Other IPS devices are part of a more sophisticated event lifecycle designed to help the security analyst not only detect the IPS alert, but also follow-through to be sure that problems are identified and resolved.

The IDS lifecycle processes of alerting, investigation, and resolution can be translated into the IPS product space as well—if this is in fact how you want to handle IPS alerts. For organizations that are looking for behavior-based and rate-based IPS built into a firewall, following through on every incident and event is probably not part of an overall security strategy. However, for organizations that maintain dedicated security staff that want to know why an IPS alert occurs —and take action based on these alerts—more sophisticated management supporting the IDS lifecycle is needed.

Management system performance is another aspect to specify carefully, particularly when the need to store events and forensics data can build up massive databases. If you plan to keep a significant amount of old data for investigative, trend matching, or regulatory reasons, you should make an effort to estimate the amount of data to help IPS vendors properly size the management console.

While forensics and alerting levy the greatest demands on IPS management systems, there are other enterprise-class management characteristics that need to be considered when defining your requirements. For example, signature-based IPS device vendors will release signature updates every few days as the threat landscape of the Internet evolves. A management system needs to support this updating in a way that meshes with your own configuration control requirements. For example, if you require that any updates to any security device be handled through a formal change control process, the management system has to support this process.

Finally, the traditional characteristics of any enterprise-class management system should be part of your evaluation criteria or requirements specification. In security devices, this often includes delegated management or role-based management (or both), reporting systems, and scalability to multiple IPS devices.

STEP 6:
Evaluate an IPS

Once you've completed all the steps in this white paper, you should actually test any IPS you're considering. At this stage of the IPS market, a test using your own network with your own traffic is the only test that will tell you whether or not the product is going to meet your requirements.

Although an IPS test doesn't require that you refine your policy completely, you should have a good idea of your network topology and security policy. Without this information, you won't be able to tell whether or not the IPS can work with your policy.

At this stage, it's also important to look at the flexibility of the IPS configuration. Can you actually express the policy you know you'll want to use in this product? For example, some IPS products don't let you easily manage exception lists for traffic that should not be inspected, or traffic that should bypass specific signatures. If you have a large and diverse network, this kind of flexibility may be important.

The recommended strategy for IPS evaluation is to put the device or combined firewall into "alert only" mode. (An IPS that doesn't have "alert only" mode should be rejected out of hand.) Rather than actually preventing intrusions, the IPS simply tells you what it would have done. When using this strategy, make sure you let the IPS run for several weeks. Until you build up a set of events, you won't know whether the product can handle the load you're going to offer it.

Once you have some confidence that the IPS isn't going to melt down your network, your evaluation should proceed to full blocking mode. When you do this, make sure you plan sufficient time each day—typically a half day, or more if your network is large or has many Internet-accessible severs—to investigate every alert, and to hunt down the false positives.[7] Even if you haven't taken the time to create a full security policy as part of your evaluation, you should be investigating most alerts. It's critical to get a feel for whether or not the IPS will actually work in your own network.

In any IPS, you should see occasional false positives. These are a natural result of an untuned system. (An IPS that does not throw any false positives ever is probably not actually working.) You should be able to fine-tune the security policy before you go into blocking mode, but still there may be false positives once you go into blocking mode. Be prepared for these, and be prepared to react quickly as they pop up. Also, remember that while clear problems will show up at your help desk in a few seconds, occasional failures may take a week or more before they begin to percolate up into support channels. When planning your testing methodology, allow for sufficient time so these "low and slow" problems will surface.

If you plan to investigate alerts, you should be sure to test the ability of the IPS to support your own "alert lifecycle." Most security managers have a specific methodology they follow to go from alert to qualification to investigation to resolution and finally to policy change. The IPS management system should support your planned methodology and style so that it is easy to handle alerts. You don't want to invest in an IPS that is hard for you to use.

With blocking enabled, it is also useful to try and 'stress test' the IPS. If you don't have commercial testing tools to inject additional load across the IPS, you can use open source tools that will increase the load of both attack and benign traffic. You may not be able to take the device to its breaking point or to precisely measure the change in behavior, but you should try to increase load by 50% or even 100% to observe the behavior of the system.

Finally, even though you may be far down an expensive evaluation cycle, it's important to step back and ask yourself whether the product you're considering and the associated capital and operational expenses give you sufficient return-on-investment for the level of security you'll be picking up. While the continuing cost of an IPS is not as high as an IDS would be, the investment in an IPS will range from simply checking a box on a firewall to enable the IPS up to installing devices and management consoles at critical points in your network. Many security professionals go down this path with an idealized idea of the value or effectiveness of IPS products. While IPS can offer significant value in improving the security posture of networks, putting that value into words just before you dive into deployment can help cement the requirements and value for IPS, as well as provide a realistic set of expectations within your organization.

[7] If an IPS sales person promises you that their product will not generate any false positives, you should thank them politely and ask them to leave.