

Title: Six Steps to Selecting the Right IPS for Your Network

By Author: [Joel Snyders](#)

Every enterprise has a firewall, but most still suffer from network security problems. IT professionals are acutely aware of the need for stronger protective technologies, and network equipment vendors are anxious to fill in the gap. Network Intrusion Prevention Systems (IPSeS) 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. The IPS market is overflowing with products with a wide spectrum of features that are suitable for a wide array of environments.

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. To decide what IPS is right for your network, follow our six-step strategy that begins by asking the question "Why am I buying an IPS?" and ends with a plan for testing an IPS in your own network.

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

The most critical step to making good decisions about security products for your network is to first know what 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 or at the core, employing signature-based technology to trap some of the bad things that pass through the network. Or, you could be more focused on mitigation of denial-of-service attacks, protecting a server farm and ensuring availability. 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.

This isn't a comprehensive list; it's just a starting point for conversations about the possible reasons to add more protective technologies to your network.

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.

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 that help segment the market into broad, overlapping categories, the underlying design goals and capabilities still vary widely.

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 Type of IPS you need.

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.

There are fundamentally three approaches in current Network IPS products: signature-based (including protocol anomaly) IPS, rate-based IPS, and behavioral IPS. (Non-Network IPS, such as wireless IPS or host-based IPS, are not part of this discussion.) While the leading products may include some pieces from all three approaches, each product has a fundamental direction it follows---signature-centric, or rate-centric, or behavior-centric – with the other two approaches being secondary.

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. 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. Most signature-based products will include some protocol anomaly measures in their repertoire.

Signature-based IPS technology, critical to catching and blocking common exploits, also 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 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. Many 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 are the most commonly used IPS technology.

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 take an active part in monitoring, controlling, and filtering 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.

Behavioral IPS tracks the flows and traffic patterns of a network. 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 more an intrusion reaction and alerting technology, and not a prevention technology.

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.

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. As 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. For standard firewalls, performance is easy to measure because metrics such as connection rate, maximum simultaneous connection count, and goodput 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 protocols 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.

Your traffic may also cause variations in performance. 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 HTTP, you would see very different performance characteristics.

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.

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.

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, 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.

Some firewalls have an "IPS function" placed into the device simply to satisfy a checklist requirement as part of a Unified Threat Management (UTM) offering. These IPSeS should be avoided, both because of their low level of threat protection, and because of their awkward and unusable management systems.

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.

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 important. IPS products vary in their management philosophy from "virtually no continuing management" to "very high management requirement" styles. 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 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: Many IPS products also have IDS capabilities, offering intensive logging, IDS signatures in addition to IPS signatures, and packet capture facilities. IPS products with this type of capability are a great addition to any network, but only if you have the staff and expertise to use them. Buying forensics capabilities that you don't or can't use is an expensive mistake.

Network visibility: Because IPSes 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 and trends at a glance.

Event alerting and correlation: SEM (security event management) tools commonly gather and correlate data from multiple sources. If you don't have a SEM, some IPS management systems have SEM capabilities.

Performance of the management system: If you plan to keep 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.

In addition to IPS-specific features, the traditional characteristics of any enterprise-class management system should be part of your evaluation criteria. This might include 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 the first five steps, you should actually test any IPS you're considering. A test using your own network and traffic is the only way to tell whether or not the product is going to meet your requirements.

You should start with 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.

Start with the IPS in "alert only" mode. Rather than actually preventing intrusions, the IPS simply tells you what it would have done. 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, 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 servers -- to investigate every alert, and to hunt down the false positives. 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 natural; 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 start blocking, but still there may be false positives once you begin. Be prepared for these, and be prepared to react quickly as they pop up. Also, remember that while some 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. Allow for sufficient time so these "low and slow" problems will surface.

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.

=====
Portions of this article originally appeared in Information Security Magazine, a TechTarget publication.