

**Buyer's Guide
For Intrusion Prevention Systems (IPS)**

Table of Contents

Introduction.....	3
Executive Summary.....	5
Quick Checklist.....	8
Detailed Buyer's Checklist.....	10

Introduction

Security has long been the number-one concern of CIOs, IT directors, and network managers. These corporate employees, responsible for the security of corporate information assets, depend on a collection of generally available security tools, including firewalls, authentication tools, VPNs and intrusion detection systems, to prevent the infiltration of their corporate networks by unauthorized individuals. However, the ever-increasing quantity and sophistication of network attacks continue to pose unforeseen threats to corporate networks, with the number of successful attacks increasing every year. The business costs of insufficient network security, resulting in the high number of security breaches, are astronomical:

- *According to the 2002 CSI/FBI survey of U.S. corporations, government agencies, financial institutions, medical institutions and universities, 90% of the respondents detected computer security breaches in the last 12 months.*
- *80% acknowledge financial loss due to computer breaches.*
- *Organizations that quantified their loss reported \$455,848,000 in financial losses due to compromises.*

Many companies believe they are secure when they implement a firewall. While firewalls are certainly the first-line of defense and an absolute requirement for any company connecting to the Internet, they cannot be the only line of defense.

For example, a firewall determines which traffic to allow or deny by applying a predefined policy that is created by the IT or network administrator. This policy is comprised of “accept” and “deny” rules for various criteria, such as the source, destination, and protocol for each communication. Most firewalls allow the set of protocols that enable organizations to do business on the Internet—such as SMTP, FTP, HTTP, and DNS—and keep out traffic that may pose a threat to internal systems.

However, the ways for attackers to misuse allowed traffic and breach the corporate network are innumerable. A quick search on Web sites that track known vulnerabilities, such as CERT (www.cert.org) and SecurityFocus (www.securityfocus.com), demonstrates the alarming reality:

Number of vulnerability notes issued for the following standard protocols:

CERT		Security Focus	
SMTP:	28	SMTP:	96
FTP:	333	FTP:	318
HTTP:	1356	HTTP:	453
DNS:	72	DNS:	73

While the specific number of vulnerabilities reported varies from site to site, the implications are the same: traffic deemed “acceptable” by a firewall can contain hidden threats that leave the corporate network vulnerable. What if, embedded in that apparently innocuous SMTP

traffic, a harmful Trojan is ready to unleash itself on your mission-critical corporate data? Firewalls were simply not designed to protect companies from all of these threats. A layered approach to network security is needed to ensure that critical assets are secure.

To provide the additional layers of security, companies must implement solutions that protect against all types of attacks embedded within network traffic.

Several approaches exist to provide this additional layer of security:

- **Add human resources** to undertake the resource-intensive task of pouring over logs, from a variety of network devices, to correlate information, identify potential attacks, and investigate each one to determine its innocuousness or success.
- **Implement simple traffic monitoring devices** that scan network traffic and alert network administrators of suspicious traffic. This approach still requires significant human resources to investigate each alert triggered by the suspicious network packets, to determine their relevance and then manually follow-up to react and recover from the intrusion.
- **Implement intrusion detection and prevention devices** that not only monitor network traffic but actually detect attacks embedded within the traffic and can drop the offending packet or connection during the detection process, so that it never reaches its intended destination and does not affect the corporate network.

With corporate intellectual property and business-critical data at stake, understanding and selecting the right approach for a company's specific requirements is a strategic imperative. This guide was designed to assist corporate decision-makers in understanding the issues involved in selection and implementation of this additional layer of security, focusing on intrusion detection and prevention options.

The following sections provide a framework for evaluating these solutions, some top level questions and a specific list of questions that evaluators can use to identify the features and functionality of each security solution to ensure the company can compare products and select the one that best meets the needs and requirements of their enterprise.

Executive Summary

Enterprises trust information security systems to protect their company from serious threats and significant financial losses. Nothing less than the very livelihood of the company is at stake—which makes the selection and implementation of security solutions that prevent unauthorized network attacks and threats a strategic requirement. This section is intended to assist decision-makers and evaluators in understanding the essential criteria to use in evaluating intrusion detection and prevention systems.

Prior to conducting a feature-by-feature comparison, decision-makers should frame their evaluation using the following six (6) criteria. An effective intrusion detection and prevention device should:

1. Provide comprehensive protection for all types of network attacks.

As with all security systems, the comprehensiveness of protection provided by an intrusion detection and prevention system is a critical element to its ability to accurately identify threats and effectively secure the network. Yet many products fall short. The inherent complexity of network traffic, which includes the vast number of protocols at both the network (IP, TCP, UDP, ICMP, etc.) and application (HTTP, FTP, SMTP, DNS, POP3, IMAP, etc.) layers, provides attackers ample vulnerabilities to exploit. Combine the inherent complexity with the fact that attacks come in different shapes and forms, and attackers have a virtual buffet to choose from when they are attacking your network. If a system does not support one of these protocols or types of attacks, it will ignore and miss the attack, leaving the enterprise network and its valuable corporate data unprotected. To thwart the attacker's efforts, a device needs to be able to address and protect against all types of traffic and potential attacks.

2. Ensure a high degree of accuracy.

Accuracy is key to an effective and efficient intrusion detection and prevention system. To create a high level of accuracy, a system must be able to track all network communications, interpret the intent of each individual communication, and then make a security decision, based on accurate evidence of an attempted attack perpetration. If the product isn't accurate, attacks may go undetected or benign traffic may be alerted on as an attack. Attacks that go by undetected is the worst scenario because it means the network is vulnerable and the administrator has to try and figure out what happened from scratch. However, if the number of false alarms that a device sends out overwhelms the administrator, or even outnumber the real ones, it is also very detrimental, wasting valuable time and resources chasing down false alarms and eroding the trust that an administrator has in the system. The system needs to be reliable and accurately detect all the attacks in the network.

3. Process traffic in a highly efficient manner.

Efficient data processing, to ensure that all functions of the system are optimally performed, is another important element of effective intrusion detection and prevention. The system must process traffic quickly, make security decisions instantly, and present that information to the network manager in a timely fashion, ensuring the administrator has a real-time view of the system at all times. A slow system that cannot keep up with the rapid flow of network traffic can mean missed attacks and an increasingly vulnerable corporate network. An administrator should also be able to choose the level of performance for the device to meet network requirements and ensure that both fast Ethernet and Gigabit network segments can be protected. The device needs to perform in an optimal manner, so that the administrator knows exactly what is going on in the network at any given time.

4. Protect against an attack without latency.

Whether an intrusion detection and prevention system can stop the attack from ever reaching its victim is the cornerstone to its prevention capabilities. How effective is an intrusion detection system that has to rely on another system to try to prevent an attack? The answer is obvious, but many intrusion detection products do just that, sending a request to a firewall or even the victims themselves to try to end the attack. All of these mechanisms come after the attack has already reached the victim, so even when successful, they require the network administrator to investigate exactly how much the attack was able to do before it was stopped. Any device that introduces latency to the prevention response, is not able to offer true prevention. A truly effective solution can actively prevent attacks during the detection process and drop the malicious traffic. This ensures it never reaches its intended victim, keeping the enterprise network and sensitive, mission-critical data safe and secure.

5. Deliver ease of use.

The ease of use of intrusion detection and prevention system translates directly into greater control and a higher degree of security. If the system enables administrators to quickly view pertinent, critical information and make adjustments, network managers can ensure the network is efficiently protected from the latest threats and the most up-to-date security policy is in force. If a device is hard to control and understand, administrators are going to waste time trying to find the information they need to do their job. It is important that the solution enables both a quick summary of the most important types of events, as well as a way to quickly drill into the raw data and manipulate it to analyze individual incidents. Providing a granular level of control in an intuitive manner to security managers, not only ensures that the system meets the company's specific security requirements, but also that valuable IT time and resources are not misspent or wasted.

6. Simple installation and maintenance.

In today's highly distributed, global enterprises, the intrusion detection and prevention product must be both easy and cost-effective to install and maintain. Companies simply cannot spare the time and resources required to update each individual device within the corporate network every time a change is made to the enterprise's security policy or a new attack is detected from which the enterprise must be protected. Quick system installation, security policy definition, and easy, global updates from a single, centralized location ensure that enterprise IT teams scattered around the globe can have a comprehensive, real-time view of the system and the network.

Quick Checklist

In the previous section, the critical criteria in which to frame the evaluation of intrusion detection and prevention systems was laid out. This section strives to provide enterprises with a quick checklist of some of the top-line questions to consider in each of the six key criteria. For more in-depth questions that enable a side-by-side comparison of different solutions, go to the Detailed Buyer's Checklist that follows this section.

1. Comprehensive Protection Against All Types of Attacks

- What methods does the product use to detect attacks?
- What network protocols can it handle? (e.g. IP, TCP, UDP, and ICMP)
- What application protocols can it handle? (e.g. HTTP, FTP, SMTP, DNS, POP3, IMAP)
- Can the product detect different stages of attack perpetration? (e.g., attack setup, attack in progress, attack compromise)
- What kinds of attacks can it detect? Does it cover:
 - Known attacks with patterns
 - Unknown attacks or attacks without patterns
 - Attacks that use interactive traffic
 - Attacks that overload a resource (DoS attacks-Syn Floods)
 - Attacks spanning multiple connections
 - Attacks using application or transport ambiguity
 - Recreational attacks
 - Layer 2 attacks
- Does the solution have different systems designed for different network segments to ensure security can be deployed throughout the network?

2. Highly Accurate Attack Detection

- How many detection mechanisms does the product use to detect attacks?
- Are these detection mechanisms working together and sharing information?
- Is the product able to track network communications in a Stateful manner?
- What mechanisms does the product have in place to interpret the data? (e.g. reassembly, de-fragmentation, normalization)
- How does the device determine where in the traffic to look for an attack?
- How does the product reduce the potential for false alarms?

3. Highly Efficient Traffic Analysis

- Are there multiple performance options to match the performance needs of different network segments?
- Are there High Availability options to reduce the chance of a single point of failure?
- What options does the system offer to deal with the traffic during the detection process?
- How does the device optimize its performance?
- Does it give the security manager a real-time view of the network?

4. Attack Prevention Without Latency

- How does the product prevent intruders from evading the system?
- What kind of response mechanisms does the product provide?
- Does the device rely on other systems to drop the attack or does it have the ability to drop the attack itself during the detection process?
- What notification mechanisms are available when an attack is detected?
- Does the product enable user-defined specific responses for individual attacks?
- How does product avoid negatively impacting important business traffic when using prevention capabilities?
- How does product avoid creating a denial-of-service situation by using prevention capabilities?

5. Ease of Use

- How does a system manager tell the product what to look for?
- Can the system look for different attacks in the same traffic and react differently?
- How do you make changes to the system's behavior, and do you need to maintain separate configurations for each device?
- How do you update the security policy?
- Is it easy to view and control the data presentment?
- Can you export logs to a SQL Database, CSV or XML files for other types of analysis?
- How does the product facilitate attack investigation?
- How do you refine the level of the data that you are looking at? (e.g. can you go from the log down to the actual packet data or up to rule that triggered it)
- Is there a way to get a quick summary view of the most important security events on your network, so you can monitor whether there are any events that need immediate attention?
- Can the summary information be manipulated?
- How easy is it to move from summary information into the particular attacks details?
- What kinds of reporting features does it offer?
- Can the system be monitored using existing enterprise monitoring tools (e.g. support SNMP-MB II)?

6. Quick Installation and Maintenance

- What is the architecture of the system?
- How is the system delivered? (e.g. appliance, software, combination of the two)
- Where are the logs collected and stored?
- Are there different options that accommodate administrator preferences for installing and configuring the system? (e.g. CLI, Web-Based)
- How do system administrators access the system's information?
- How many policies are required to manage a distributed network?
- How do you change a policy or update a signature for all of the sensors?
- Is the system able to easily integrate into the network (e.g. support VLANs)?
- Can a single device protect multiple network segments?

Detailed Buyer's Checklist

Evaluation Date: _____

Evaluated By: _____

FEATURE	NetScreen-IDP	Vendor # 2	NOTES
<i>Comprehensive Attack Protection</i>			
The number of methods the device can use to detect attacks	8		
Product can detect known attacks with patterns (code red)	Stateful Signatures, Protocol Anomaly		
Product can detect unknown attacks and attacks that cannot be characterized (buffer overflows, DNS cache poisoning, future sendmail exploits)	Protocol Anomaly Detection, Backdoor Detection		
Product can detect unauthorized interactive traffic (Trojans, worms, back orifice)	Backdoor Detection		
Product can detect reconnaissance attacks that span multiple connections (port scans, network scans)	Traffic Anomaly Detection		
Product can detect attacks designed to overload a resource (DoS attacks-Syn Floods)	DoS Detection, Protocol Anomaly Detection		
Product can detect attacks at layer 2 (ARP spoofing)	Layer 2 Detection		
Product can detect IP Spoofing	IP Spoof Detection		
Product can imitate services and lure attackers to attack those non-existent services to reduce the noise of recreational attackers (script kiddies)	Network Honeypot		
The number of network protocols the product supports: <ul style="list-style-type: none"> • TCP • IP • UDP • ICMP • ARP 	Yes Yes Yes Yes Yes		

deployments to ensure each network segment can be protected	the product line		
Solution for large central site	Yes		
Solution for medium central site/large branch office	Yes		
Solution for small remote office	Yes		

Accurate Attack Detection			
The product uses multiple attack detection mechanisms to minimize the likelihood that attacks can go by undetected	Yes		
The product's detection mechanisms look at the network traffic independently	No		This has the potential to create false alarms
The product's attack detection mechanisms work together and share information	Yes		
The product is capable of using the most appropriate method to detect each type of attack	Yes		
The product can apply state information to its signature detection	Yes		Can only be done if the mechanisms are sharing information
The product can pinpoint where to look for an attack pattern match (signature), by: <ul style="list-style-type: none"> • Connection type (any, client to server, server to client) • Flow (control, auxiliary, both) • Fixed offset • Packet • Stream • Line • Service Field 	Yes Yes Yes Yes Yes Yes		*Most systems can only look at the packet level
Number of service fields where attacks are identified	225+		
The product normalizes network traffic to interpret the data exactly as the destination machine would interpret it	Yes		
The product reassembles packets for accurate reconstruction of data streams	Yes		
The product has an open signature format	Yes		
The product allows the user to create custom signatures using a GUI	Yes		
The product allows the user to create custom signatures using a script	Yes		
The product enables users to define Stateful Signatures	Yes		

scanner, can be configured to focus on the incidents that pose the greatest risk to the network			
Attacks are grouped to make it easy to identify the most critical threats to the network	Yes		
Attacks can be grouped according to user-defined custom groups	Yes		

Efficient Traffic Processing			
Product can participate in a Gigabit environment (Fiber Gigabit Ethernet Standard)	Yes		
Product can achieve full line-speed to meet the performance requirements of most network connections (Fast Ethernet)	Yes		
Product can scale down to a smaller network segment	Yes		Burst rates at full-line speed can be achieved for less than a second
Products can be clustered to increase performance capacity	Yes		
Product offers standalone HA to ensure there is no single point of failure	Yes		
Product offers status monitoring of all devices in HA Clusters	Yes		
Product offers load balancing/sharing	Yes		
Product has a fail open option	Yes, with a Bypass unit		Need to purchase Bypass unit
Product performs log indexing for fast log presentment and look up	Yes		
Product is able to do parallel signature pattern matching, rather than succession matching	Yes		
Product only looks for signature patterns in the relevant portions of traffic, from the stream level down to the service field, where the attack can be perpetrated	Yes		
Product reacts to attacks as soon as it is detected, without any latency	Yes		
Product presents the system information in a timely manner for real-time analysis	Yes		
Product offers system status monitoring	Yes		

Attack Prevention			
The product can be configured to react to detected attacks differently	Yes		
The number of response options the product offers	7		
The product relies on another device to try to stop the attack (firewall, client, server, etc.)	No, not unless configured to do so		Firewall Signaling/Blocking, TCP Resets
The product can drop the attack itself, without latency	Yes		Must be in-line to do this
The product can drop the malicious packet	Yes		
The product can drop the malicious connection	Yes		
The product can close the connection (send TCP Reset to both the client and server)	Yes		
The product can close the connection on the client (TCP Reset)	Yes		
The product can close the connection on the server (TCP Reset)	Yes		
The product can reconfigure the firewall to block IP address of the attacker	Yes		Not recommended, due to DoS potential
The product can do nothing to the connection	Yes		
The number of notification/logging options the product offers	7		
The product enables the user to configure exactly how many packets before and after the attack it should capture	Yes		
The product can archive logs, along with related packet data, for later analysis	Yes		
The product can export logs into a CSV file for analysis by third party tools	Yes		
The product can be configured to send an e-mail	Yes		
The product can be configured to send e-mails to different addresses based on the type of attack	Yes		
The product enables user-defined responses for individual attacks	Yes		

Ease of Use			
The product uses a rulebase to define behavior at a granular level	Yes		
The product enables the user to define rules that dictate: <ul style="list-style-type: none"> • exactly what traffic the product should look for • what attacks in that traffic to look for • what to when the attack is identified • which sensor(s) to apply the rule 	Yes Yes Yes Yes		
The product can easily be configured to look for different things in the same traffic and react differently	Yes		
The user has to make an all or nothing decision about whether or not to look for an attack (turn the signature on or off)	No		Point-to-point management
The product ships with template rulebases to help the user get started	Yes		
The product can be tuned through a single enterprise-wide policy, dictating different behavior for each enforcement point (sensor)	Yes		Centralized, rule-based management
When changes are made to the policy and pushed live, the product automatically updates each individual sensor with the appropriate changes	Yes		
Each enforcement point needs its own individual policy that must be touched every time there is a change or update	No		Point-to-point management
When a sensor is added to the system, a policy push will automatically apply all appropriate rules in the policy to it	Yes		
The user can assign flags to highlight specific attacks or events	Yes		
The product comes with severity guidance for attacks	Yes		
The product enables the user to override the severity levels within the rule base for particular attack	Yes		
The product allows the user to filter the information in the logs	Yes		
The product enables the user to have multiple filtered views	Yes		
The product enables the user to have multiple views (policy, multiple filtered log views, system status) up at the	Yes		

same time in different windows to facilitate investigation			
The product enables users to easily clear filters or save them in their preferences for future use	Yes		
The product enables users to add comments to the rules and logs, facilitating communication among the security team	Yes		
The user can easily move from the log up to the rule/security policy that triggered it and down to the packet data to determine exactly what happened	Yes		Sometimes called Closed Loop Investigation
The product provides a summary of each log, highlighting the pertinent information of the packet	Yes		
The product enables users to export log data for later investigation: SQL database CSV file XML file	Yes Yes Yes		
The vendor provides updates to the product (e.g. new signatures) on a regular basis to protect against the latest threats	Yes, with support contract		
The user can automatically apply the updates provided by the vendor to the system	Yes		
The user can pick and choose individual signatures within a signature update.	Yes		
When updates to the system are done, the product will alert the user to potential overrides of custom signatures to ensure this is the intent of the user	Yes		
The product offers information about the attack from the interface	Yes		
Product provides quick links to additional attack definitions for further assessment	Yes		
Product enables the user to turn off information if they so desire	Yes		
Product offers a comprehensive help section	Yes		
The number of pre-defined reports the product offers	12		

The product enables users to export the report data for easy dissemination	Yes		
The product offers multi-level forensic investigation	Yes		
The product provides a dashboard summary view of important events to quickly understand what is going on in the network	Yes		
The product enables the user to define what they want to see in the dashboard	Yes		
The product enables the user to drill down from the summary information to the specific data about the event they need	Yes		
The product allows the user to retrieve log and packet information directly from the reports	Yes		
The product enables the user to visually correlate sources against hosts, against attacks	Yes		
The product can be monitored with any SNMP-based monitoring product, which the enterprise may already be using to monitor other devices	Yes		

Installation and Maintenance			
<p>The device uses a three-tier architecture, consisting of:</p> <ul style="list-style-type: none"> • Sensor (enforcement point, detects and prevents attacks) • Centralized management server (collects all logs, stores the policy, configuration and user information) • Distributed graphical interface (the way the user accesses and interacts with the system) 	<p>Yes</p> <p>Yes</p> <p>Yes</p>		
<p>The product's architecture makes it easy to scale for a large distributed enterprise</p>	<p>Yes</p>		
<p>The product offers multiple interfaces for installation and configuration, so the user can determine which they prefer to use</p>	<p>Yes</p>		
<p>The product offers a simple, Web-based interface for installation, configuration and changing settings</p>	<p>Yes</p>		
<p>The product offers a text browser/CLI for installation, configuration and changing settings</p>	<p>Yes</p>		
<p>All communications between the tiers of the product are encrypted and authenticated</p>	<p>Yes (RSA and Blowfish)</p>		
<p>The product can be remotely accessed in a secure manner by an unlimited number of remote users</p>	<p>Yes</p>		
<p>The product is delivered as an appliance for easy installation</p>	<p>Yes</p>		
<p>The product can manage a distributed enterprise using a single policy</p>	<p>Yes</p>		
<p>All updates made by the user are automatically updated by the product to the sensors in a secure manner</p>	<p>Yes</p>		
<p>The product can integrate into the network</p>	<p>Yes</p>		
<p>The product supports virtual routers to allow traffic to be forwarded to specific virtual or physical interfaces</p>	<p>Yes</p>		
<p>The product can protect multiple segments</p>	<p>Yes (up to 20 forwarding interfaces per sensor)</p>		
<p>The product supports 802.1Q VLAN tagging on any interface</p>	<p>Yes</p>		