An ISS Technical White Paper


# Intrusion Protection—The Latest Weapon in the Network Security Arsenal

6303 Barfield Road • Atlanta, GA 30328 Tel: 404.236.2600 • Fax: 404.236.2626

## Executive Summary

No one today doubts the need for secure online business operations. The challenge is how to do so efficiently and effectively. Which is why intrusion protection has become an essential element of the layered approach to enterprise-wide network security, which also includes firewalls, virtual private networks, anti-virus scanning and more. And like many other aspects of network security, intrusion defenses are undergoing a transformation from passive detection to active, in-line protection against attacks.

The network intrusion detection system (NIDS) has done a very effective job of performing the tasks it was designed to do: monitor host or network traffic, detect attacks or suspicious behavior and alert administrators. As a result, this technology is now widely accepted and well understood, and has become a standard practice for organizations of all sizes.

The active, in-line network intrusion protection system, or network IPS, is the natural evolution of NIDS. After all, effective intrusion protection requires the best possible intrusion detection. The network IPS is a powerful and logic extension of basic NIDS technologies that builds on this foundation to provide significantly improvements in both the effectiveness and cost-effectiveness of network security. When deployed as part of pervasive solution that unites network, server and desktop protection across fixed, wireless and remote networks, the network IPS helps lower the total cost of ownership (TCO) and improve the return on investment in enterprise-wide security.

The network IPS has three distinct advantages over a purely NIDS-based solution:

- *Accuracy* – With centralized management and automated correlation, the network IPS easily handles the massive amounts of constantly changing security events and incidents without overlooking potentially critical details.
- *Performance* – Automation of the entire security cycle minimizes the time required to close the loop between the discovery of a new vulnerability or recognition of an attack, and the ability to act or defend against it.
- *Lower Costs* – The network IPS significantly lowers TCO by using advanced data collection, correlation and analysis to minimize false alarms or false positives, limits the need for manual intervention in the security process, and automates the discovery and repair of potential security exposures.

The Internet Security Systems approach is different from any other. Internet Security Systems extends the IPS vision across NIDS, vulnerability assessment, and comprehensive security research, analysis and services. This Dynamic Threat Protection™ process moves enterprise protection from semi-automated intrusion detection and response via semi-compatible point products to an automated, intelligent security *platform* that extends across networks, servers and desktops, and includes fixed, wireless and remote systems.

This white paper addresses the need for and requirements of next-generation network intrusion protection systems. The material is organized into three main sections followed by a brief conclusion. The first section highlights the evolution of network protection, culminating in the present-day use of network and host intrusion detection systems. The next section focuses on the recent evolution from passive *detection* to active network intrusion *protection* systems with the advent of the network IPS. The final section provides an introduction to the RealSecure® Guard network IPS from Internet Security Systems. Although some content is technical, most of the material is suitable for an executive-level management audience.

## The Evolution of Network Protection

Once upon a time, enterprise networks were totally secure. Organizations employed private lines to link the headquarters with remote offices. These links were expensive, but they were also virtually impenetrable. The advent of higher speed modems and dial-up access systems brought about the need for simple password-based security provisions. Yet the many benefits of greater access continued to far outweigh the modest increase in risks for most organizations.

Then along came the Internet and everything changed. As the first truly global data network, the Internet created possibilities that were unimaginable with the inherently secure Public Switched Telephone Network. Higher throughput. Ubiquitous reach. Any-to-any connectivity. Unprecedented flexibility. And all at a much lower cost. But there was one huge problem: The Internet was designed as an "open" network with no security. The open approach helped make the Internet the world-wide success it is today, of course, but its use in commercial applications quickly led to the rise of a whole new industry in network security.

Internet security began with the firewall. Just the name itself—"firewall"—made the IT staff feel safe and secure behind it. Analogous to locks on doors and windows in the home, firewalls create barriers to the corporate network's entry points. But the cyberspace of the Internet soon brought about another change: cybercrime. Hackers, virtual vandals, disgruntled employees and other social misfits soon found they could exploit the firewall's inherent limitations for personal gain or revenge, or just the notoriety of being able outsmart the security experts. For this reason, firewalls are no longer enough to protect enterprise networks.

The state-of-the-art in today's hacker technique, the blended or hybrid attack, is designed specifically to exploit the weaknesses that exist in most network security provisions. Take as examples the recent Code Red virus, Nimda worm and Trojan horse attacks that were cleverly perpetrated to bypass traditional firewall and anti-virus defenses. Distributed denial of service attacks take a similar approach: Break up an assault into smaller elements that go undetected individually, then reunite once through the perimeter to launch the attack. The only way to stop this latest form of attack is with a comprehensive network protection solution.

### Traditional Network Intrusion Detection Solutions

The objective of network protection is to keep the bad guys out, while continuing to let the good guys in. Traditionally this objective has been achieved by deploying numerous network intrusion detection system (NIDS) agents throughout the enterprise network. Each NIDS monitors its assigned network segment and reports a series of security events or incidents back to a centralized management console.

Commonly employed as another layer of network defenses, the NIDS complements the firewall "door and window locks" with the equivalent of "security cameras" that monitor network traffic, detect attacks and intrusions, then alert network administrators to take corrective or pre-emptive action. The key to effective intrusion detection is a thorough and accurate analysis of the events reported by NIDS agents deployed throughout the enterprise. Such an analysis must employ a number of different and effective intrusion detection techniques. Simply put: A single technique is incapable of detecting all possible intrusions. Indeed, accurately identifying attacks requires a number of advanced IDS techniques employed combination of methods must be employed.

The more powerful detection techniques include full stateful 7-Layer protocol analysis, advanced pattern-matching, baseline anomaly analysis, and heuristics and behavior analysis. Protocol analysis is a particularly effective technique that requires the full reassembly and decoding of all packets in a network connection or session. The analysis normally performs a validation of the protocol's RFC compliance, including the utilization of all fields in the protocol's header. A

combination of these and other techniques has allowed the NIDS to become a well understood and widely accepted weapon in the network security arsenal.

Thus far, intrusion detection systems, in concert with firewalls, anti-virus scanners and other defenses, have done a fairly admirable job of mitigating the risks of using the Internet for most organizations. With the more recent hybrid or blended attacks designed to circumvent these defenses, however, the results have been mixed. With Nimda, for example, the better NIDS solutions sounded the alarm. But in many organizations, Nimda was able to penetrate the perimeter and attack vulnerable servers. In effect, the NIDS warning simply did not come in time to permit an effective response.

Which is why a growing number of organizations want more. They recognize that to achieve the vision of dynamic protection, which allows for an instantaneous response to changing vulnerabilities and threat conditions, they need a more active intrusion protection system: one that not only detects the attacks, but also intervenes immediately and automatically to block the attack. What these organizations need is the next step in network security: intrusion *protection*.

## From Intrusion Detection to Intrusion Protection

At its core, the NIDS is a passive device; it can only detect an attack in progress. Responding to any threat detected requires human intervention. For this reason, a growing number of organizations are now augmenting or even replacing passive intrusion *detection* systems with active, in-line intrusion *protection* systems. Instead of detecting an attack and issuing an alert that requires operator intervention, the network intrusion protection system, or network IPS, intervenes to block the attack. As a real-time, in-line access control system that forms the last line of defense, the advent of the network IPS signals the next advance in the state-of-the-art for intrusion protection solutions.

Of course, detection remains a necessary element of active, in-line intrusion protection. For without first accurately detecting attempted attacks and intrusions, it is impossible to prevent these threats and, ultimately, protect the network's attached assets. The network IPS must, therefore, build on the solid foundation established by the NIDS to enable security personnel to become more proactive in defending the enterprise against the dynamic threat spectrum, as well as to provide significantly improved value for the security investment already made in NIDS technology. With the network IPS, companies get a highly automated, intelligent security platform with tightly integrated operations that lead to much faster response times, focus administrator attention on areas needing the most attention, and anticipate security exposures before they become business liabilities.

Naturally, as the next new thing in network security, the network IPS lends itself to hyperbole in the marketplace. Some of the claims are incredible—quite literally—while others may seem more reasonable. Because the hype currently far exceeds the reality, the usual advice applies: *caveat emptor*. Some vendors claim their network IPS products can detect and block every single known and any future unknown attack, stopping it dead in its tracks, without introducing any latency or creating a performance bottleneck, while security administrators sit back and enjoy their new-found freedoms. But with such a "plug-and-play" IPS, the equivalent of the inevitable false positive is far more devastating than a mere waste of a security administrator's time. Blocking *legitimate* traffic "automatically and immediately" undermines the very purpose of a business presence on the Internet.

### Network IPS Requirements

To be truly effective, the network IPS must be part of well-designed and well-managed intrusion protection solution. It bears repeating: In-line protection requires detection, and for detection to be accurate against hybrid or blended threats, it must also be pervasive. Which is why savvy security

professionals make certain they have intrusion detection working properly *before* deploying intrusion protection systems.

The network IPS is typically deployed as an additional layer of network security to protect a mission-critical server or network segment, which may contain a server farm or legacy host system. It can also be deployed as a complete protection solution replacing a traditional NIDS. In either role, the network IPS provides both a static blocking capability, similar to a firewall, and a dynamic blocking capability, which prevents exploits of vulnerabilities specific to the server or segment being protected. Static blocking, for example, may be employed as a preemptive measure to block access to the unused server processes that are turned on automatically by the operating system, and to restrict access to certain other processes to specific IP addresses. Dynamic blocking then builds on this foundation as a real-time response to a current attack on any available process. The specific mechanisms used by a network IPS vary, but normally involve limiting or terminating sessions, dropping packets, or stripping packets of their harmful content—a process called "neutering".

Because the network IPS performs multiple functions, some solutions are distributed among multiple devices. For example, the IPS may be designed to take control of a separate, third-party firewall. Such implementations can be quite complicated and difficult to manage, however. Which is why the better network IPS solutions are designed as standalone, full-featured systems.

A few other features are just as important in a robust network IPS solution. The first is reliability. Because a network IPS must sit in-line, it becomes a single point of failure. The system should, therefore, have built-in redundancies and, ideally, a by-pass capability that lets all traffic through in the event of unrecoverable failure. As an in-line system, throughput and latency become important considerations. Achieving high performance is particularly challenging because the best network IPS solutions perform a full stateful 7-Layer protocol analysis, which is the only meaningful way to detect the traffic anomalies that are symptomatic of a first-strike attack.

Perhaps the most important element of effective network intrusion protection is centralized management. In fact, a centralized management console is the only way to detect hybrid or blended attacks, manage other network, server and desktop protection solutions, and assure enterprise-wide conformance with the security policy. Without this important tool, security administrators constantly face trying to make sense of the relentless onslaught of events from agents located throughout the enterprise. Manually monitoring a large network to detect and respond to threatening activity can be a difficult, time-consuming and error-prone task. This problem caused first-generation intrusion detection systems to report an excessive number of "false positives." A false positive is an alert regarding a potential threat that has little or no actual adverse consequences. For example, a false positive occurs when an attack is launched on Microsoft Windows NT servers at a data center that has only Linux servers. Of course, one company's legitimate security threat may be another's false positive, and getting too few false positives may indicate that alert thresholds are configured too high. But too many false alarms can become like the boy who cried "Wolf!"

To minimize false positives and false alarms to focus available resources exclusively on real threats, the centralized management console requires an automated correlation system. Correlation is the process of distilling raw threat event data into prioritized, actionable information. The raw data is reported as security events by all NIDS, network IPS and other agents deployed throughout the enterprise on network segments, desktops and servers, and potentially from firewalls, ant-virus scanners and other security systems.

Accurate correlation of these security events is vitally important to detecting and preventing attacks, determining the actual business impact of an attack, and monitoring and reducing overall organizational risk, while significantly reducing false positives and false alarms. Simply put: Correlation is the best means available for prioritizing the seemingly endless number of security events into those that pose a real threat and, therefore, warrant immediate attention. Indeed,

pervasive intrusion protection with automated correlation is what makes it possible to move from having monthly or quarterly vulnerability assessment snapshots to obtaining a daily or even a real-time dynamic awareness of security risks on an enterprise-wide basis.

With automated correlation, intrusion protection solutions provide an effective way to secure large networks from both known and first-strike attacks, and also allow smaller organizations to achieve more sophistication without a staff of security experts. When properly configured, correlation eliminates the drudgery of merely trying to find threats and, instead, focuses available resources on taking responsive and/or pre-emptive measures to deal with those threats.

## The RealSecure Guard Network Intrusion Protection System

Leveraging its best-of-breed network intrusion detection technologies, Internet Security Systems introduced RealSecure Guard as the industry's first active, in-line intrusion protection system that secures networks, including mission-critical segments and systems, by actively blocking attacks. Guard provides its robust intrusion protection in a single, centrally managed device that combines firewall services with intrusion detection and dynamic blocking to thwart attacks in real-time. As traffic passes through the Guard network IPS, it is either blocked by static firewall rules, or analyzed in-depth using sophisticated protocol analysis techniques, then blocked dynamically based on evidence of intrusion, attack or misuse. Once improper behavior is detected, Guard blocks the attack by dropping or neutering packets.

Unlike many other in-line intrusion protection systems, the industrial strength packet processing drivers in Guard ensure accurate, high-speed monitoring without compromising performance. With Guard, Internet Security Systems offers another layer of network protection to augment the market-leading RealSecure network intrusion detection systems. Guard also extends the RealSecure umbrella of protection to mainframe and mid-range hosts by monitoring all traffic to and from these mission-critical systems in real-time and without affecting their operation.

**How Guard Works**

Guard has three interfaces: two that sit in-line for ingress and egress of network traffic, and one for secure out-of-band management. As packets are reassembled and decoded, then analyzed in a virtual TCP/IP protocol stack, Guard takes one of three possible actions:

- If the traffic is permitted, Guard discards the analyzed copies as the original packets pass intact.
- If the packets contain an attack that is not immediately threatening, Guard reports the event to the *RealSecure*® *SiteProtector*™ management console and logs the traffic to a forensic evidence file for later analysis.
- If the attack poses an immediate threat, Guard actively blocks the traffic by dropping the packets, then reports and logs the event.

The ability to configure Guard systems for different responses depending on the severity of the attack allows intrusion protection to be integrated into the overall security strategy at a comfortable pace.

RealSecure Guard offers a number of state-of-the-art features that make active in-line intrusion protection the next logical step for organizations with NIDS-based solutions, or a practical first step for organizations just now considering an intrusion protection solution. Here is just a partial list of the powerful features available with the Real Secure Guard network IPS:

- *Stateful Protocol Analysis* (see sidebar) – Guard uses a powerful combination of sophisticated, 7-Layer protocol analysis and attack pattern matching to fully and accurately interpret network activity. To detect traffic anomalies that signal a potential first-strike attack, Guard is able to decode and analyze some 60 different protocols. These techniques are the most accurate and modern methods of intrusion detection available, and are designed specifically to foil hacker evasion techniques. The combination allows Guard to detect both known and previously unknown attacks, making it immune to attempts to evade purely signature-based or pattern matching systems.
- *Real-world Accuracy* – Network IPS solutions that strictly enforce "RFC Compliance" can produce false positives that block legitimate traffic, causing expensive troubleshooting delays and downtime of mission-critical servers. Guard's comprehensive protocol analysis techniques

are sophisticated enough to account for the common protocol violations found in popular applications to prevent such false positives.

- *Bi-directional Protection* – Guard analyzes traffic in both directions to detect not only inbound attacks, but also to detect misuse and potential outbound threats that might be part of a hijacking or distributed denial of service attack. For this reason, Guard spots and stops hostile traffic from either entering or exiting a protected segment, while allowing legitimate traffic to pass unhindered.
- *Integral Firewall* – Guard has a built-in firewall that is reconfigured dynamically and automatically to block malicious traffic. The firewall can also be configured with static blocking rule sets based on ports or offending IP addresses for preemptive protection. While not intended as a replacement for separate, full-featured firewalls, the integral firewall in Guard can eliminate the need for a separate firewall in some deployment scenarios.
- *Built-in Reliability* – In the event of catastrophic failure or power loss, Guard has an "off-line" bypass feature that continues to pass traffic. This standard feature can also be used to take Guard out of service for diagnostic or other purposes with absolutely no disruption to normal traffic.
- *Line-rate Performance* – Guard can accurately and effectively identify and block attacks on a fully-utilized 100 Mbps full-duplex connection. Guard also handles traffic encrypted at the Secure Sockets Layer (SSL) at the full line rate, and its architecture is extensible to Gigabit-per-second data rates.
- *Centralized Management* – Centralized command and control in RealSecure SiteProtector integrates and correlates events for seamless, enterprise-wide management of a pervasive RealSecure Protection Platform solution.
- *Ease of Use* – Guard is a single, self-contained system that requires no external or third party devices to deploy intrusion protection for a mission-critical network segment or server. Such simplicity is ideal for large-scale environments or for any organization with limited resources.

**Stateful Protocol Analysis**

The stateful 7-Layer protocol analysis capability in Guard delivers the most sophisticated intrusion protection available with these advanced features:

- Both packet- and session-level analysis with full protocol decode and reassembly of all packets involved in a network connection or session
- Bi-directional analysis of both in-bound and out-bound traffic to detect misuse, hijackings and distributed denial of service attacks
- Attack fingerprint and signature databases, updated regularly, to protect against all known attacks
- Protocol and baseline anomaly detection to identify potential first-strike attacks
- Protocol field pattern matching and validation, complete with real-world RFC compliance testing, are also used to detect first strike attacks
- Full line-rate analysis of SSL-encrypted traffic, which is increasingly common in high security eBusiness applications and transactions
- Customizability and user-specified rule sets that allow network defenses to be tailored to unique or unusual needs

## Pervasive and Dynamic Protection with the RealSecure Protection Platform

RealSecure Guard is part of the *RealSecure Protection Platform*, the industry's most complete and comprehensive intrusion detection/protection solution. The RealSecure Protection Platform is a highly integrated, centrally managed best-of-breed protection solution that encompasses security assessment, intrusion detection. protection and response, policy enforcement and decision support. With its pervasive coverage of networks, desktops and servers, the RealSecure Protection Platform is purpose-built to afford the industry's highest levels of protection against both attacks and misuse.

One reason for the success of the RealSecure Protection Platform is recognition by Internet Security Systems that intrusion protection defenses are somewhat different than other network security provisions. Like most network threats, intrusion techniques are subject to constant change. It seems as soon as one form of attack becomes known, another one emerges. But unlike other threats, intrusions can be distributed or segmented in clever ways to escape detection by some defenses. For this reason, pervasive coverage and centralized coordination of

enterprise-wide intrusion protection defenses are essential to thwarting both known and first-strike attacks.

As the industry's premiere choice in intrusion protection, the RealSecure Protection Platform offers the most pervasive coverage with a choice of network, desktop and server systems. The desktop and server agents add an important layer of protection to NIDS and network IPS agents because they detect and protect against attacks at their intended targets. Network-based devices cannot confirm the success or failure of an attack on a server or host, especially when the attack originates from behind the perimeter of defense—that is, from the supposed "trusted" side of the network or directly from the console. Server-based solutions can also help lower the total cost of ownership by being deployed where network-attached devices are either impractical or too costly.

The RealSecure Protection Platform gives enterprises and service providers alike pervasive protection with more choices and options than any other solution available. Products and services in the RealSecure Protection Platform include:

- The *RealSecure SiteProtector* centralized command and control console that provides unified configuration and data analysis capabilities for RealSecure network, server and desktop protection solutions. SiteProtector simplifies large-scale RealSecure Protection Platform deployments through cost-efficient, central management and monitoring that reduce administrative demands on staff and other operational resources. The ability to deploy SiteProtector in a tiered arrangement allows for virtually unlimited scalability and also facilities greater flexibility when using managed services options. The *RealSecure SiteProtector security fusion module* adds advanced data correlation and analysis techniques to rapidly and automatically determine the likelihood of a successful attack based on vulnerability assessment information. Fusion dramatically lowers the occurrence of false positives, which helps focus available resources more productively on mitigating real threats.
- RealSecure Network 10/100 and Network Gigabit agents are policy enforcement network intrusion detection systems that leverage Internet Security Systems' advanced NIDS capabilities to provide sophisticated application-level protocol analysis and pattern-based detection technologies for 10/100 Mbps and Gbps network segments. These agents complement Guard intrusion protection systems by enhancing detection capabilities through pervasive, enterprise-wide coverage.
- The *RealSecure Server* host intrusion detection system (HIDS) defends servers and applications through sophisticated distributed firewall capabilities, server-based intrusion detection and response, security policy distribution, and malicious code elimination. All server and application protection operations are tightly integrated with each other, and report to the SiteProtector console for simplified configuration and management.
- The *RealSecure Desktop* protection agent for corporate workstations employs a defense-in-depth methodology with both an integral personal firewall and advanced protocol analysis for robust intrusion protection. The behavior-oriented protocol analysis engine fully decodes and structurally assesses entire network transmissions to block all forms of potential intrusion, including fragmented attacks. The RealSecure Desktop agent installs "silently" on desktop and laptop PCs, and integrates seamlessly with most third-party VPN and anti-virus products.
- RealSecure assessment solutions provide comprehensive network and server vulnerability assessments for measuring online security risks and ensuring security policy compliance. Vulnerability assessment management serves as the very foundation of a pervasive intrusion protection solution and is essential to accurate event correlation.
- The Internet Security Systems *X-Force*™ organization is the foundation for all RealSecure solutions. This cutting-edge R&D team actively turns security research into product improvements, allowing Internet Security Systems customers to respond far more rapidly and effectively to threats. Leveraging this security knowledge, the Internet Security Systems comprehensive suite of X-Force Protection Services provides comprehensive customer support with industry-leading professional security services, educational services, emergency response, and 24/7 remotely managed security services and technical support.

## Conclusion

Ironclad security for today's Internet-accessible networks and corporate assets requires an arsenal of security provisions that includes firewalls, anti-virus scanners, VPNs, intrusion protection and more. Protecting corporate assets from attack and misuse, whether from internal or external threats, is now more critical than ever with the proliferation of new hybrid threats, such as Nimda. Traditional intrusion detection systems have now been enhanced to include centralized management, automated correlation, stateful protocol analysis and other features that make the NIDS more effective than ever before. But some organizations want more. They want active intrusion protection systems capable of securing mission-critical resources in real-time.

Internet Security Systems continues its tradition of industry leadership with support for these and other state-of-the-art capabilities. Indeed, ISS remains the only vendor offering a total and pervasive intrusion protection solution complete with vulnerability assessment tools, NIDS and network IPS solutions, desktop and server protection systems, centralized management with automated correlation, managed and professional services, and more.

The migration from detection to protection has always been central to the Internet Security Systems vision of network defenses that are integrated and automated to achieve the vision of full Dynamic Threat Protection with its immediate response the ever-changing spectrum of potential threats. RealSecure Guard has made taking this important step from detection to protection a practical and affordable one today for the enterprise and service providers alike.

To learn more about how your organization can benefit from the industry's #1 choice in intrusion protection, call ISS at 888-901-7477 or visit www.iss.net.

# # #

**About Internet Security Systems (ISS)**

Founded in 1994, Internet Security Systems (ISS) (Nasdaq: ISSX) is a pioneer and world leader in software and services that protect corporate and personal information from an ever-changing spectrum of online threats and misuse. Internet Security Systems is headquartered in Atlanta, GA, with additional operations throughout the Americas, Asia, Australia, Europe and the Middle East. For more information, visit the Internet Security Systems Web site at www.iss.net or call 888-901-7477.