McAfee®

Next Generation Intrusion Detection Systems (IDS)

Table of Contents

I.   Introduction

Today, the network is the business. Driven by business needs, enterprises and government agencies have developed sophisticated, complex information networks, incorporating technologies as diverse as distributed data storage systems; encryption techniques; Voice over IP (VoIP); remote and wireless access; and Web services. These networks have become more permeable as business partners access services via extranets; customers interact with the network through e-commerce transactions or Customer Relationship Management (CRM) processes; and employees tap into company systems through Virtual Private Networks (VPN).

For hackers, these well-traveled paths make networks more vulnerable than ever before and—with relatively little expertise—hackers have significantly impacted the networks of leading brands or government agencies. Cyber crime is also no longer the prerogative of lone hackers or random attackers. Today disgruntled employees, unethical corporations, even terrorist organizations all look to the Internet as a portal to gather sensitive data and instigate economic and political disruption.

With networks more vulnerable and hackers equipped to cause havoc, it's no surprise that network attacks are on the rise.

According to a 2001 report by Computer Security Institute (CSI) and the FBI, 70 percent of respondents acknowledged that their networks were attacked over the previous twelve months and 30 percent didn't know whether their networks were attacked and couldn't be sure! In addition, Denial of Service attacks increased by an astonishing 33 percent over the same period. And all this took place across networks, where firewalls had been installed in 90 percent of instances.

It's clear that enterprises and government agencies need security vendors to step up and deliver innovative solutions that effectively protect their networks from malicious attacks and misuse.

Being introduced is the industry's first realtime network intrusion prevention platform that takes Intrusion Detection Systems (IDS) to a new level. This integrated hardware and software platform is based on the company's McAfee® IntruShield® architecture, which delivers comprehensive protection from known, first strike (unknown), and Denial of Service (DoS) attacks—at multi-gigabit speeds. IntruShield also provides the flexible management capabilities needed to administer the varied security policies required for the individual departments, diverse geographies, and separate functions that make up global businesses and vital government agencies.

This paper discusses the benefits of IDS products and the technologies that a next generation IDS must provide to deliver effective protection and presents the capabilities of the IntruShield architecture.

II.   The Need for IDS

When most people think of network security, they think *Firewall*. Firewalls are widely deployed as a first level of protection in a multi-layer security architecture, primarily acting as an access control device by permitting specific protocols (such as HTTP, DNS, SMTP) to pass between a set of source and destination addresses. Integral to access policy enforcement, firewalls usually inspect data-packet headers to make traffic-flow decisions. In general, they do not inspect the entire content of the packet and can't detect or thwart malicious code embedded within normal traffic. It should be noted that routers also offer some rudimentary protection through packet-filtering processes.

While firewalls and router-based packet filtering are necessary components of an overall network security topology, they are insufficient on their own.

Network IDS products inspect the entire content of every packet traversing the network to detect malicious activity. This content inspection technique provides deeper packet analysis compared to a firewall or a router. Intrusion Detection Systems are effective when sophisticated attacks are embedded in familiar protocols, such as an HTTP session, which would normally pass undetected by a firewall. It's not surprising that the processing power required for an Intrusion Detection System is an order of magnitude higher, when compared to a firewall product.

Permeable modern networks have made IDS products essential tools as security engineers strive to detect, analyze, and protect networks against malicious attack. As a result, IDS products are being deployed outside and inside firewalls and are quickly becoming mainstays in *best practice* secure network implementations.

III.   Understanding IDS

IDS products can be split into broad categories—Host IDS and Network IDS products. Host IDS products protect an end system or network application by auditing system and event logs. Network IDS products can be deployed on the network, monitoring network traffic for attacks. A Network IDS can sit outside the firewall, on the demilitarized zone (DMZ), or anywhere inside the private network.

Typically Network IDS products focus their efforts around one of three areas—Signature Detection, Anomaly Detection, or Denial of Service (DoS) Detection.

Signature Detection

Protecting against known threats.

Hackers often attack networks through tried and tested methods from previously successful assaults. These attacks have been analyzed by network security vendors and a detailed profile, or attack signature, has been created. Signature detection techniques identify network assaults by looking for the attack *fingerprint* within network traffic and matching against an internal database of known threats. Once an attack signature is identified, the security system delivers an attack response, in most cases a simple alarm or alert. Success in preventing these attacks depends on an up-to-the-minute database of attack signatures, compiled from previous strikes. The drawback to systems that rely mainly, or only, on signature detection is clear: they can only detect attacks for which there is a released signature. If signature detection techniques are employed in isolation to protect networks, infrastructure remains vulnerable to any variants of known signatures, first-strike attacks, and Denial of Service attacks.

Anomaly Detection

Protecting against first strike or unknown threats.

Anomaly detection techniques are required when hackers discover new security weaknesses and rush to exploit the new vulnerability. When this happens there are no existing attack signatures. The *Code Red* virus is an example of a new attack, or first strike, which could not be detected through an available signature. In order to identify these first strikes, IDS products can use anomaly detection techniques, where network traffic is compared against a baseline to identify abnormal—and potentially harmful—behavior. These anomaly techniques are looking for statistical abnor-malities in the data traffic, as well as protocol ambiguities and atypical application activity. Today's IDS products do not generally provide enough specific anomaly information to prevent sophisticated attacks and if used in isolation, anomaly detection techniques can miss attacks that are only identifiable through signature detection.

Denial of Service (DoS) Detection

Protecting against network and system overload.

The objective of DoS and Distributed DoS attacks is to deny legitimate users access to critical network services. Hackers achieve this by launching attacks that consume excessive network bandwidth or host processing cycles or other network infrastructure resources. DoS attacks have caused some of the world's biggest brands to disappoint customers and investors as Web sites became inaccessible to customers, partners, and users—sometimes for up to twenty-four hours. IDS products often compare current traffic behavior with acceptable normal behavior to detect DoS attacks, where normal traffic is characterized by a set of pre-programmed thresholds. This can lead to false alarms or attacks being missed because the attack traffic is below the configured threshold.

No single technique or technology is the "magic bullet" to guarantee protection against current or future attacks.

In order to robustly protect enterprise and government networks against the complete spectrum of threats and vulnerabilities, all three methodologies must be employed—Signature Detection, Anomaly Detection, and Denial of Service Detection and Prevention.

Also, a next generation IDS must do more than detect attacks: it should enable accurate detection to prevent attacks from reaching and damaging critical network resources and data. Without this range of detection methods—and the performance to accurately prevent attacks—many IDS products are no more than a digital Maginot Line: while they may offer the illusion of protection, when real attacks come, defenses can be circumvented or overrun.

IV. IDS Challenges Today

Most of today's IDS products are focused on Signature Detection and are designed for sub-100Mbps shared media network environments, employing detection capabilities introduced three to four years ago. IDS products have failed to keep up with the rapid advancement in switching and bandwidth growth and the increased sophistication of attacks—as well as their sheer volume. Current IDS products often operate in a monitoring-only mode, *sniffers*, which can detect attacks but cannot effectively and reliably block malicious traffic before the damage is done.

Network security managers deploying IDS products today face a number of challenges:

Incomplete attack coverage: IDS products typically focus on Signature, Anomaly, or Denial of Service

detection. Network security managers have to purchase and integrate point solutions from separate vendors or leave networks vulnerable to attack.

Inaccurate detection: IDS products' detection capabilities can be characterized in terms of accuracy and specificity. Accuracy is often measured in *true detection rate*—sometimes referred to as the *false-negative rate*—and the *false-positive rate*. The true detection rate specifies how successful a system is in detecting attacks when they happen. The false-positive rate tells us the likelihood that a system will misidentify benign activity as attacks. Specificity is a measure of how much detailed information about an attack is discovered when it is detected. IDS products today are lacking in both accuracy and specificity and generate too many *false-positives*, alerting security engineers of attacks, when nothing malicious is taking place. In some cases, IDS products have delivered tens of thousands of *false-positive* alerts a day. There is nothing more corrosive to network vigilance than a jumpy security system, which is continually issuing false alarms.

Detection, not prevention: Systems concentrate on attack detection. Preventing attacks is a reactive activity, often too late to thwart the intrusion.

Designed primarily for sub-100Mb/s networks: Solutions have simply not kept up with the speed and sophistication of network infrastructure and cannot accurately monitor higher-speed or switched networks.

Performance challenged: Software applications running on general purpose PC/server hardware do not have the processing power required to perform thorough analysis. These underpowered products result in inaccurate detection and packet dropping, even on low bandwidth networks.

Lack of high-availability deployment: Single port products are not able to monitor asymmetric traffic flows. Also, with networks becoming a primary mechanism to interact with customers and partners, forward-thinking organizations have developed back-up systems should their current infrastructure fail in any way. The inability of current IDS products to cope with server failovers renders them virtually useless for any mission-critical network deployment.

Poor scalability: Primarily designed for low-end deployments, today's IDS products do not scale for medium and large enterprise or government networks. Here monitored bandwidth, the number of network segments monitored, the number of sensors needed, alarm rates, and the geographical spread of the network exceed system limits.

No multiple policy enforcement: Current products generally support the selection of only one security policy for the entire system, even though the product may monitor traffic belonging to multiple administrative domains—in an enterprise this could be the finance, marketing, or HR functions. This *one size fits all* approach is no longer acceptable for organizations that require different security policies for each function, business unit, or geography.

Require significant IT resources: IDS products today require substantial hands-on management—for example, the simple task of frequent signature updates can take up a lot of time and skilled engineering resources, delivering a very high total cost of ownership.

In response to these limitations, a new architecture that detects and prevents known, unknown, and Denial of Service attacks was developed for even the most demanding enterprise and government networks. The remainder of this paper will discuss the innovative technologies and capabilities of the IntruShield architecture.

V. Introducing McAfee IntruShield Security Architecture

The IntruShield architecture delivers realtime network intrusion prevention at multi-gigabit speeds. Comprehensive protection is delivered on an integrated, purpose-built platform that can scale across highly-available networks. The schematic below describes how IntruShield delivers the industry's most robust Intrusion Detection System and redefines the network IDS space.


IntruShield Architecture
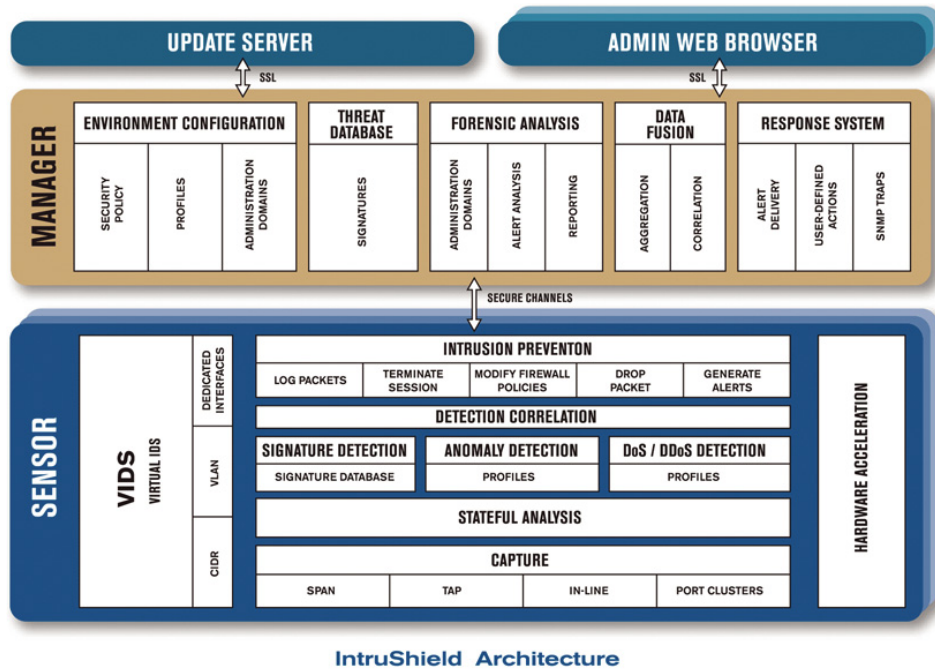
The IntruShield architecture consists of three major components: sensor system, management software, and the update server. Let's look more closely at how the IntruShield architecture enables new functionality to be delivered in each of these components.

Capture

The IntruShield architecture enables sensor systems to capture network attacks in a number of ways:

Switched Port Analyzer (SPAN) and Hub Monitoring: Hub ports or SPAN ports from one or more network switches can be connected to the IntruShield system's detection ports. Response actions such as resetting a TCP connection can often be injected by the sensors using the same port.

Tap Mode: Network communication is monitored in both directions of a full duplex Ethernet network link. By fully capturing all of the traffic on a link, a clearer understanding of the source and nature of the network attack can be delivered—and can provide the detailed information needed to thwart future attacks. This full-duplex monitoring capability allows IntruShield systems to maintain complete state information. Response actions include firewall reconfiguration or initiating a TCP reset through dedicated response ports.

In-Line Mode: IntruShield systems sit in the data path, with active traffic passing through them. The IntruShield system prevents network attacks by dropping malicious traffic in real time. Preventative action can be custom-defined at a highly granular level, including automated dropping of DoS traffic intended for a specific Web server. Wire speed prevention and highly available operation enable IntruShield system deployment in mission-critical environments.

*Port Clustering* allows traffic monitored by multiple ports on a single IntruShield system to be *aggregated* into one traffic stream for state and intrusion analysis. This feature is especially useful in environments with asymmetric routing, where request and response packets may traverse separate links. A single IntruShield system can monitor multiple links and maintain accurate and complete state information.

A single IntruShield system with multiple interfaces can offer comprehensive Perimeter Protection by monitoring all segments connected to the firewall in either full-duplex tap mode or in-line mode.

The IntruShield architecture also enables IDS systems to become integral to High Availability topologies (active-active or active-passive) and asymmetrically routed deployments. Here IntruShield systems can fail over to a hot standby.

VI.   Stateful Analysis

With the IntruShield architecture delivering a rich set of capture capabilities, let's look at the in-depth and intelligent analysis of the captured data. IntruShield systems enable for stateful analysis of network traffic and packet reassembly.

IntruShield systems maintain complete state information—inspecting the entire content of the data packet—as it traverses the monitored network links. State information is captured and updated in real time. Maintaining state information enables sensors to gain context for attack detection, delivering higher accuracy of attack detection.

IP Defragmentation and TCP Stream Reassembly

Within this analysis IntruShield performs full IP defragmentation and TCP stream reassembly, emulating the traffic received by the end-systems being protected. This is important because hackers utilize attack techniques that fragment malicious code across multiple data packets and often reorder these packets to further evade detection. Once these packets reach their target the host reassembles the data and the malicious code does its damage. IntruShield reassembles these packets before they hit the intended target, providing new levels of protection.

Protocol Analysis

The IntruShield architecture enables detailed analysis of all major protocols, ensuring highly accurate attack detection rates. In addition to leveraging protocol analysis for buffer overflow detection—a major class of recent attacks—protocol parameters are also made available to write powerful and accurate user-defined signatures.

Traffic Normalization

In addition, the IntruShield architecture's Traffic Normalization functionality—available when the system is operating in in-line mode—removes any traffic protocol ambiguities; meaning that the traffic being interpreted by IntruShield systems and the traffic received at the protected end-system are identical. IntruShield systems remove any traffic protocol ambiguities, protecting the end systems by cleaning up potentially harmful traffic in real time. Traffic normalization thwarts any attempts to evade the Intrusion Detection System while boosting attack detection accuracy. While operating in tap mode, IntruShield systems issue alerts when uncovering protocol ambiguities.

VII.   Signature, Anomaly, and Denial of Service Detection

IntruShield's architecture builds on its thorough attack analysis methodologies by adding the industry's most comprehensive Signature, Anomaly, and Denial of

Service detection techniques. This section looks in detail at how the IntruShield platform approaches each detection method.

1. Signature Detection

For Signature Detection, the IntruShield architecture's innovative and patented technology combines a Stateful Signature Detection Engine, a sophisticated Signature Specification Language, *User-Defined Signatures*, and Realtime Signature Updates.

The architecture outlined below enables the delivery and the maintenance of the industry's most complete and up-to-date Attack Signature Database.

Signature Specification Language

IntruShield's architecture is bolstered by a proprietary high-level Signature Specification Language. IntruShield decouples signatures from application software and, in this unique architecture, signatures simply become table entries, which can be updated in real time using an intuitive user interface and used immediately by the Signature Engine.

Today's IDS products often deliver new signatures via a software *patch*; slow to develop because it must be quality-assured against the entire IDS software application, and cumbersome to install because the system must be re-booted. On the other hand, IntruShield ensures that high-quality new signatures can be deployed quickly—without requiring system reset—by separating signatures from the sensor software. Decoupling signatures from sensor application code also allows signature writers to focus on the *art* of signature writing, rather than worry about building their signatures into an updated application patch.

Stateful Signature Detection Engine

The IntruShield architecture's Signature Detection Engine employs powerful context-sensitive detection techniques that leverage state information within data packets, utilize multiple token matches, and hunt down attack signatures that span packet boundaries or are in out-of-order packet stream.

User-Defined Signatures

IntruShield empowers network security engineers to write custom signatures through an innovative Graphical User Interface (GUI) that can utilize specific fields and data obtained through the system's protocol analysis capabilities, or state information gathered via IntruShield's Analysis mechanisms.

Realtime Signature Updates

The IntruShield architecture powers management software that offers an innovative realtime Signature Update process, where new signatures made available by an Update Server can be pushed out across the network in a policy-controlled automated fashion, ensuring that networks are protected as soon as a new signature is created. The IntruShield architecture also allows the network engineers to make the decision on when and if to deploy the new signature across their networks. Because the IntruShield system does not require any hardware reset or reboot to leverage new signatures, they automatically kick-in in real time.

2. Anomaly Detection

The IntruShield architecture's comprehensive Signature Detection processes are complemented by a set of Anomaly Detection techniques that allow network engineers to thwart emerging threats or first-strike attacks and create a far-reaching set of Anomaly profiles, protecting the network against current threats and future attacks.

The IntruShield architecture delivers the industry's most advanced and complete anomaly detection methodology—encompassing statistical, protocol, and application anomaly detection techniques. Example categories of anomaly/unknown attacks are new worms, intentionally stealthy assaults, and variants of existing attacks in new environments. Anomaly detection techniques can also help in thwarting denial of service attacks, where changes in service quality can be observed; and distributed DoS attacks, where traffic pattern changes (such as TCP control packet statistics) can be used by the IntruShield system to determine whether a data deluge is on the way. We'll discuss Denial of Service attacks more in the following section.

Other areas that the IntruShield architecture's anomaly detection techniques help guard against are buffer overflow attacks; backdoor malicious attacks installed via a Trojan or by an insider; stealthy scanning attacks that use low frequency, multiple launch points on the network and deliver normal looking packets; and insider violation of security policies, such as installing a game server or a music archive on the network.

3. Denial of Service Detection

The third pillar in IntruShield's detection architecture is its sophisticated Denial of Service protection technologies.

Self-Learning and Threshold-Based Detection

The IntruShield architecture employs a combination of threshold-based detection and patented self-learning profile-based detection techniques that delivers intelligence to Denial of Service detection. With threshold-based detection, network security managers can utilize pre-programmed limits on data traffic to ensure servers will not become unavailable due to overload.

Meanwhile, self-learning methodologies enable the IntruShield architecture to study the patterns of network usage and traffic, understanding the wide variety of lawful, though unusual, usage patterns that may take place during legitimate network operations.

The combination of the two yields the highest accuracy of detection for a full spectrum of DoS attacks— including distributed Denial of Service attacks, when hundreds or even thousands of servers are co-opted by a malicious programmer to strike against an enterprise or government network.

IntruShield's accurate DoS detection techniques are important because popular Web sites and networks do experience legitimate—and sometimes unexpected—

traffic surges for a particularly compelling new program, service, or application.

VIII. Detection Correlation

As we have seen, the IntruShield architecture enables numerous modes of operation that allow the system to capture malicious traffic; provides thorough attack analysis methodologies; and implements a complete set of intelligent Signature Detection, Anomaly Detection, and Denial of Service protection techniques.
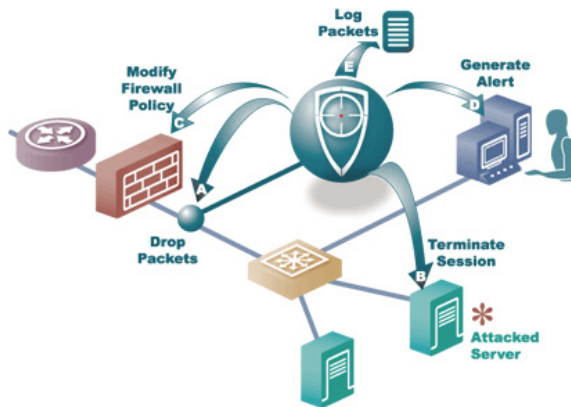
The IntruShield architecture's Detection Correlation layer connects the system's Signature, Anomaly, and Denial of Service detection functionality—and this interdependence and cross-checking of suspicious traffic yield highly accurate attack detection.

A single IntruShield system—providing comprehensive protection by monitoring public, private, and DMZ segments of the firewall—can offer correlation among these segments to yield an accurate picture of network attacks that were either blocked by the firewall or made it into the private network.

IX. Intrusion Prevention

IntruShield's architecture delivers the industry's most accurate attack detection capabilities, forming the basis for the system's attack response mechanisms. An IDS without adequate response capacity is of limited utility to network security managers. Modern IDS products must detect attacks and provide the means to deflect and stop malicious traffic.

The IntruShield architecture supplies network security managers with a full spectrum of manual and automatic response actions that can form the basis of an enterprise's or government agency's information technology security policies.

Upon detecting an attack, the IntruShield architecture enables the system to:

A. Drop Packets—Because the IntruShield architecture allows the IDS to work in in-line mode, it is able to drop or block a single packet, single session, or traffic flow between the attack source and destination in realtime, thwarting an attack in progress without affecting any other traffic.

B. Terminate Session—The IntruShield architecture allows for the initiation of TCP resets to targeted systems, attackers, or both. The network security engineer can configure reset packets to be sent to the source and/or destination IP address.

C. Modify Firewall Policies—The IntruShield architecture allows users to reconfigure network firewalls as an attack occurs by temporarily changing the user-specified access control policy while alerting the security manager.

D. Genertate Alerts—The IntruShield architecture enables an Alert Filter that allows network security engineers to sift out alerts based on the source or the destination of the security event. For example, if the IT department executes vulnerability scans from one of its own IP addresses, events originating from that address can be filtered out.

E. Log Packets—Systems based on the IntruShield architecture capture and log packets prior, during, or subsequent to the attack and can redirect traffic to a spare system port for detailed forensic analysis. This packet information acts as a record of the actual flow

of traffic that triggered the attack. When the data is viewed it is converted to libpcap format for presentation. Tools like Ethereal, a network protocol analyzer for UNIX and Windows, can be used to examine the packet log data for more detailed analysis of the detected event.

The IntruShield architecture's response mechanisms will provide the basis for the product platforms that security managers need to develop a system of actions, alerts, and logs that provide optimum protection for complex contemporary networks.

X. Virtual IDS (VIDS)

In recognition of the complexity of today's networks, IntruShield's architecture allows for the creation of multiple Virtual Intrusion Detection Systems (VIDS™). Virtual IDS domains can be set up for specific departments, geographic locations, or functions within an organization. Security policies can then be set for each Virtual IDS, providing the flexibility modern organizations need when managing a diverse set of network users.

The IntruShield architecture's Virtual IDS functionality can be implemented in three ways. Firstly, by attributing Virtual Local Area Network (VLAN) tag(s) to a set of network resources; secondly, by protecting a block of IP addresses utilizing Classless Inter-Domain Routing (CIDR) tags; and thirdly, by dedicating IntruShield system interfaces to protect the network resources in particular department, geography, or organizational function.

CIDR-based VIDS implementation allows granularity down to an individual host level with /32 mask. For example, DoS attacks can be identified and responded to with unique policies for individual hosts.

XI. Hardware Acceleration

The powerful functionality of the IntruShield architecture—from Capture through to the implementation of Virtual IDS—is made possible by dedicated, purpose-built, proprietary hardware that provides the performance required to accurately detect and then prevent network intrusions at wire-speed without packet loss. To be effective, IDS products must have at least an order of magnitude higher performance than even the most cutting-edge firewall systems.

Almost every task undertaken by IntruShield systems benefits from hardware acceleration. For example, IntruShield's Signature Processing capabilities require hardware to accelerate repetitive signature detection tasks, such as string matches. As a result the IntruShield architecture can support thousands of attack signatures at multi-gigabit data rates—and at the same time continue to detect and prevent first strike and Denial of Service assaults.

It's clear that today's IDS products, many of which rely on software running on generic PC or server hardware, are not up to the task of detecting and preventing network intrusions.

The system architecture outlined above delivers the core technology needed by network security engineers to protect network resources. As can be seen from the architecture diagram, the management and control features enabled by IntruShield allow engineers to impact and administer policies for all areas of attack detection and prevention.

XII. Management and Control

The IntruShield architecture makes management and control functions available via a secure, Web-based, highly-graphical administration tool. Role-based access controls allow network professionals within an organization to be granted different access privileges. The powerful, yet easy-to-use graphical tools provide scalable

management—from performing routine administrative and reporting tasks to establishing network-wide security policies and administrative domains.

Environment Configuration allows network engineers to control security policies across the network, access environment profiles to understand the technology mix at a particular location, and set the best security regime for each administrative domain. Users can also control and add to the constantly updated Threat Database, where the latest attack signatures sit, along with those from previous known attacks as well as other attack profiles. Thanks to IntruShield's architecture, engineers are able to perform Forensic Analysis by logging traffic and examining alerts. Attacks can be dissected for hard information on their source, capabilities, and intended target—and new signatures can be created from this valuable data. In addition, the IntruShield architecture's Data Fusion functionality aggregates and correlates information from a wide variety of network assaults to provide managers with meaningful attack summaries, discerning the broad patterns of attack activity and providing a basis for an organization's security policies. And lastly, IntruShield enables network engineers to have full control over a comprehensive Response System, allowing detailed management of alerts, user-defined actions, SNMP traps, and attack reports.

Completing the IntruShield architecture is the Update Server. The Update Server, a high-availability server, automatically pushes out new signatures and system updates to the management systems at customer locations. Signature and system updates are then deployed within customer networks by security professionals using advanced tools provided within the IntruShield Management platform.

XIII. Summary

As we have seen, the threats against enterprise and government networks are real and growing. IDS products have been gaining ground as an appropriate response to known, first-strike, and denial of service attacks and have become an important tool complementing network firewalls. However, the permeable nature of complex and sophisticated networks—that bring customers, employees, business partners, and the general public in direct contact with network resources—has made networks more vulnerable to

attack. The current generation of IDS products does not provide the breadth of detection techniques, or the accuracy and performance, to prevent network attacks from reaching their intended targets.

XIV. About the Author

Dr. Fengmin Gong is the Chief Scientist for the McAfee Network Security Technologies Group, where he is responsible for driving the continued innovation of IntruShield's security architecture—leveraging his expertise in areas such as signature, anomaly, and denial of service detection. Before his work on IntruShield, Dr. Gong was Director of Advanced Networking Research at MCNC, a provider of sophisticated electronic and information technologies and services aimed at businesses and government agencies.

While at MCNC and earlier at Washington University, he was involved in advanced security and networking projects for agencies such as DARPA, NSA, NSF, NLM, and NASA. During his time at MCNC he was also Adjunct Assistant Professor of Computer Science at North Carolina State University.

In a distinguished academic and research career, Dr. Gong has written and contributed to nearly forty research papers on network intrusion, anomaly detection, secure collaboration, multi-media content delivery, and network quality of service. Dr. Gong has presented his research at industry events such as IEEE technical forums, as well as SIGGRAPH, DISCEX, NOMS, and ISCEX.

XV. About McAfee Network Protection Solutions

McAfee Network Protection Solutions keep both large and smaller distributed networks up and protected from attacks. Best-of-breed network protection solutions in the portfolio include McAfee Security Forensics, providing data capture, terabyte storage capacity, and information retrieval to analyze important security events and network threats; and McAfee IntruShield, delivering network-based intrusion prevention.

McAfee IntruShield

IntruShield, a part of the McAfee Network Protection Solutions family of products, is a unique cutting-edge technology that prevents intrusions *on the wire* before they hit critical systems. Highly automated and easily managed, IntruShield is designed with such flexibility that it can be implemented in a phased approach—that overcomes the false-positives inherent with today's legacy intrusion detection systems—and thus enables you to develop the right policy for blocking in your unique IT infrastructure. For example, you can deploy in-line to notify and block known attacks, and to notify-only on unknown attacks.

Or you can implement complete blocking, but just for business-critical network segments. IntruShield is delivered in a high-speed appliance which is able to scan traffic and assess threat levels with blinding speed, even on gigabit networks. It can be used at the edge or in front of key *core* resources. IntruShield has been crafted to satisfy both the security and network administrators as it stops a wide range of network attacks but does so with network latencies typically less than 10 milliseconds. IntruShield also looks for anomalous behavior and includes specialized analysis to find new denial of service *mass attacks*.

XVI. About McAfee

With headquarters in Santa Clara, California, McAfee, Inc. (NYSE: MFE) creates best-of-breed computer security solutions that prevent intrusions on networks and protect computer systems from the next generation of blended attacks and threats. Offering two families of products—McAfee System Protection Solutions, securing desktops and servers, and McAfee Network Protection Solutions, ensuring the protection and performance of the corporate network—McAfee offers computer security to large enterprises, governments, small and medium sized businesses, and consumers. These two product portfolios incorporate McAfee's leading product lines. For more information, McAfee can be reached at 888-847-8766 or on the Internet at http://www.mcafee.com/.

McAfee, Inc. 3965 Freedom Circle, Santa Clara, CA 95054, 888.847.8766

McAfee® products denote years of experience and commitment to customer satisfaction. The McAfee PrimeSupport® team of responsive, highly skilled support techni-cians provides tailored solutions, delivering detailed technical assistance in managing the success of mission-critical projects—all with service levels to meet the needs of every customer organization. McAfee Research, a world leader in information systems and security research, continues to spearhead innovation in the development and refinement of all our technologies.

www.mcafee.com