Security Monitoring Realities and Futures

A White Paper from Vigilinx

Vigilinx
Digital Security Solutions

Table of Contents

Any inquiries should be addressed to:
info@vigilinx.com

www.Vigilinx.com

Executive Summary

Security Monitoring

Realities and Futures

The network security landscape has experienced a sea change over the past five years. The introduction of Intrusion Detection Systems (IDS) has altered the way organizations protect themselves. An IDS improves visibility into the inner workings of the network, identifying security issues in a way that was previously impossible.

However, an organization implementing IDS faces a number of significant challenges. An IDS does not operate in the same manner as more traditional security controls such as firewalls. Many inexperienced implementers discover that the technology can overload their monitoring capabilities, rendering them useless. More importantly, because IDS is in its infancy, these problems will only increase unless organizations move quickly to understand them.

This paper discusses some of the major issues in implementing an IDS and strategies for achieving success. It does so within the larger context of network security protection. This approach is important because security technology is changing rapidly. The fusion of security information systems that monitor firewalls, IDS, and virus scanning will vastly increase both the capability and workload of the security operations group. By developing a workable approach to IDS now, most organizations can set the stage for integrating these newer systems into their environment in the future.

Introduction

The last decade has seen a dramatic shift in the way information technology operates. This shift has most affected the networking of workstations and servers. With the introduction of Transmission Control Protocol/Internet Protocol (TCP/IP), we have linked local systems, local-area networks, and wide-area networks with a unified protocol. As a result, all users can communicate with all other users, regardless of location.

Unfortunately, unifying the protocol has a major drawback: because every single system is now addressable, susceptibility to unauthorized access and vandalism from remote parties now exists. Network security has become the single most critical factor in maintaining the integrity of business operations. As security vendors have struggled with this problem, they have developed a series of technologies to address it, including:

- Packet filtering
- Application proxy firewalls
- Stateful packet inspection
- Virtual private networks
- Intrusion detection systems

These and other technologies address a growing list of perceived threats from the Internet and other sources. In the beginning, the firewall was intended to block traffic known to affect vulnerable internal resources. However, the recent shift to detection technology signals a fundamentally new approach. By detecting complex traffic patterns deep within an enterprise network, IDS technology enables better performance and greater functionality while enhancing overall security.

The dark side of IDS is its proclivity to generate enormous quantities of detection data. While it is imperative that organizations monitor all security technology, the problem of managing the volume of information — yet to find a complete solution — is driving two reactions:

1. Many early buyers of IDS technology have decided the struggle is not worth it. They desperately need its features, but they do not have the staff or expertise to put it to use. As a result, they have abandoned its use or greatly scaled back plans to deploy it.
2. Several service providers have begun to develop and offer managed IDS services. They are betting on their ability to decipher IDS data through intelligent placement of sensors, filtering of data, and planned information integration systems.

This paper explores some of the techniques of these service providers and looks to the future to understand the next round of security challenges.

IDS Fundamentals

Any company whose network interfaces with external networks (the Internet, customers, vendors, and partners) requires "security in depth." This means policies, controls, and technical solutions like firewalls, access control lists (ACL), and intrusion detection. Intrusion Detection Systems can serve many purposes. They allow one to validate the configuration of a firewall and they can easily show how many people are attempting an intrusion. They also introduce multiple "layers of protection" from intrusion.

As with most new technologies, intrusion detection has its share of kinks. These range from too many false detections to difficulties functioning in switched and high-speed environments. Careful deployment can overcome these problems through good understanding of IDS capabilities, which makes IDS an extremely useful tool for timely and efficient monitoring and response.

Several critical factors lead to successful IDS implementation, including:

- Understanding the current state of IDS technology;
- Determining the optimal locations to deploy IDS sensors;
- Configuring IDS filtering to reduce false positives;
- Maintaining the IDS technology through implementation of new features and system upgrades; and
- Allocating appropriate staff to implement and monitor the systems, perhaps outsourcing some or all of these duties.

Let's explore each of these factors.

Understanding the Current State of IDS Technology

Because it is new, IDS technology has seen a great deal of innovation over the past four years. Early on there were only a few dozen attack signatures, but this has grown to more than 500 today. Where there was once just one vendor, we now have more than a dozen in the marketplace. Demand has driven this explosive growth. However, while a number of published articles have focused on IDS, there still exists only limited understanding of its position in the security solutions arena.

IDS vendors have focused on the number of available attack signatures to claim competitive differentiation, and competitive analysis generally reports on how many different types of attacks a product can detect. Unfortunately, this distinction is immaterial, since the majority of implemented signatures are so outdated that few systems remain vulnerable to these attacks.

Another downside of the "battle of the numbers" is that the quality of new signatures is extremely poor. This has increased the number of false positives, which has lowered the overall level of IDS usability.

Understanding IDS shortcomings enables an organization to successfully deploy intrusion detection. While it is a valuable tool, IDS doesn't address all security issues. Rather, it functions much like a burglar alarm. As such, it cannot catch information theft committed by insiders. Moreover, its effectiveness depends upon the level and quality of monitoring it gets.

When IDS is used on a medium to large scale, it can trigger a large number of events. These are not all false positives. Many alerts are valid, indicating events like malfunctioning network equipment, system anomalies, or application errors. If a company has one or two people "in charge" of intrusion detection efforts, those people may find they have to deal with hundreds, or even thousands, of alerts on a daily basis. Unless they develop an effective strategy for doing so, they will likely discontinue their monitoring efforts.

Therefore, it's critical to determine how to address this volume on the procedural, not just the technological, level.

Determining the Optimal Locations to Deploy IDS Sensors

Traditionally, organizations begin using IDS at their perimeter. Most IDS professionals, however, recommend placement much further within the network, where damage will have the greatest impact. Nonetheless, due to the organizational dynamics of corporate environments, it is often easier to start with Internet-facing networks.

Let's look at a typical network configuration and determine the optimal IDS strategy for it. Figure 1 depicts a common network scenario, viewed at the perimeter with the Internet. The network consists of three basic layers:

1. An external facing demilitarized zone (DMZ);

2. An application gateway internal DMZ; and

3. The internal network.

The external DMZ provides the portal for the outside world to access the Internet services of this organization. Here we typically find applications such as WWW (HTTP), domain name system (DNS), and e-mail. Some companies place (anonymous) file transfer protocol (FTP) servers in this DMZ as well.

Generally, an external firewall limits services to those necessary for public access. Firewall administration is straightforward. Since all traffic types are known, the firewall should block everything else. Security experts call this type of rule, "Deny all not specifically permitted."

The internal DMZ affords a buffer between the external DMZ servers and the internal network. It functions as a "pass through" between the two environments. Communication is on a strict IP/port to IP/port basis. For example, the external WWW server may receive HTTP traffic from the Internet. Data processing necessary to complete the web transaction occurs on systems in the internal DMZ fed from customer records on the internal network. The two hosts communicate through appropriate IP/port combinations. This configuration eliminates the possibility of direct public access to sensitive data on the internal systems by maintaining a "safe zone" on the internal DMZ.

There are a total of five IDS sensors in the figure. All are configured with two network interface cards (NIC), so they can operate "invisibly." The red dotted lines in the figure represent these "stealth" interfaces, with a 0.0.0.0 IP address. The red dashed lines represent the internal, private address space. This configuration makes the IDS sensors invisible to all, even internal, users. The entire IDS monitoring solution operates over its own subnet.

With this configuration, the IDS sensors will continue to see and report on network traffic, even in the presence of a DoS attack. Another advantage is that, since the sensors and the IDS consoles are invisible, they cannot be targeted for attack. This is an important feature because hackers often receive insider assistance. If a motivated attacker knows that IDS is in place, he or she will try to flood or otherwise disable it. Once IDS is gone, the organization is blind to anything that happens on the network.

Using five sensors may seem like the ideal IDS configuration, but for most organizations it is not. One reason is that this configuration may not be appropriately matched to the way the organization operates. For example, consider the following scenarios:

- Company A is a law firm. In the external DMZ , it has a webserver and an e-mail server. The internal DMZ transfers the mail. All users are on the internal network.
- Company B is an online retailer. The external DMZ houses a webserver, e-mail server, and shopping cart server. The internal DMZ has a database server that processes the online orders.

How would the two configurations differ?

Company A definitely needs to protect the internal network, where sensitive client information is stored. Thus, IDS No. 3 should be in place. Internet availability is not a critical requirement for most similar organizations, so placing an IDS in either DMZ (IDS No. 4 and IDS No. 5) isn't a priority. However, IDS No. 2 provides extra protection, and this setup will give them very few false positives.

Company B will also require IDS No. 3, since the internal network always needs protection. Because the database server on the internal DMZ stores credit card and other crucial information, IDS No. 5 is appropriate. Further, Internet availability is critical, so they need to protect the customer-facing environment, requiring IDS No. 4.

This simple example demonstrates how, in spite of identical network architectures, different companies have different needs for IDS. More importantly, neither company requires IDS sensor No. 1, even though this is the most common placement today. By putting the sensor outside the external firewall, the organizations will observe a large number of events. Unless these data are used for trending purposes, this IDS is unnecessary and will generate false positives.

Figure 1: Typical IDS Deployment Scenario

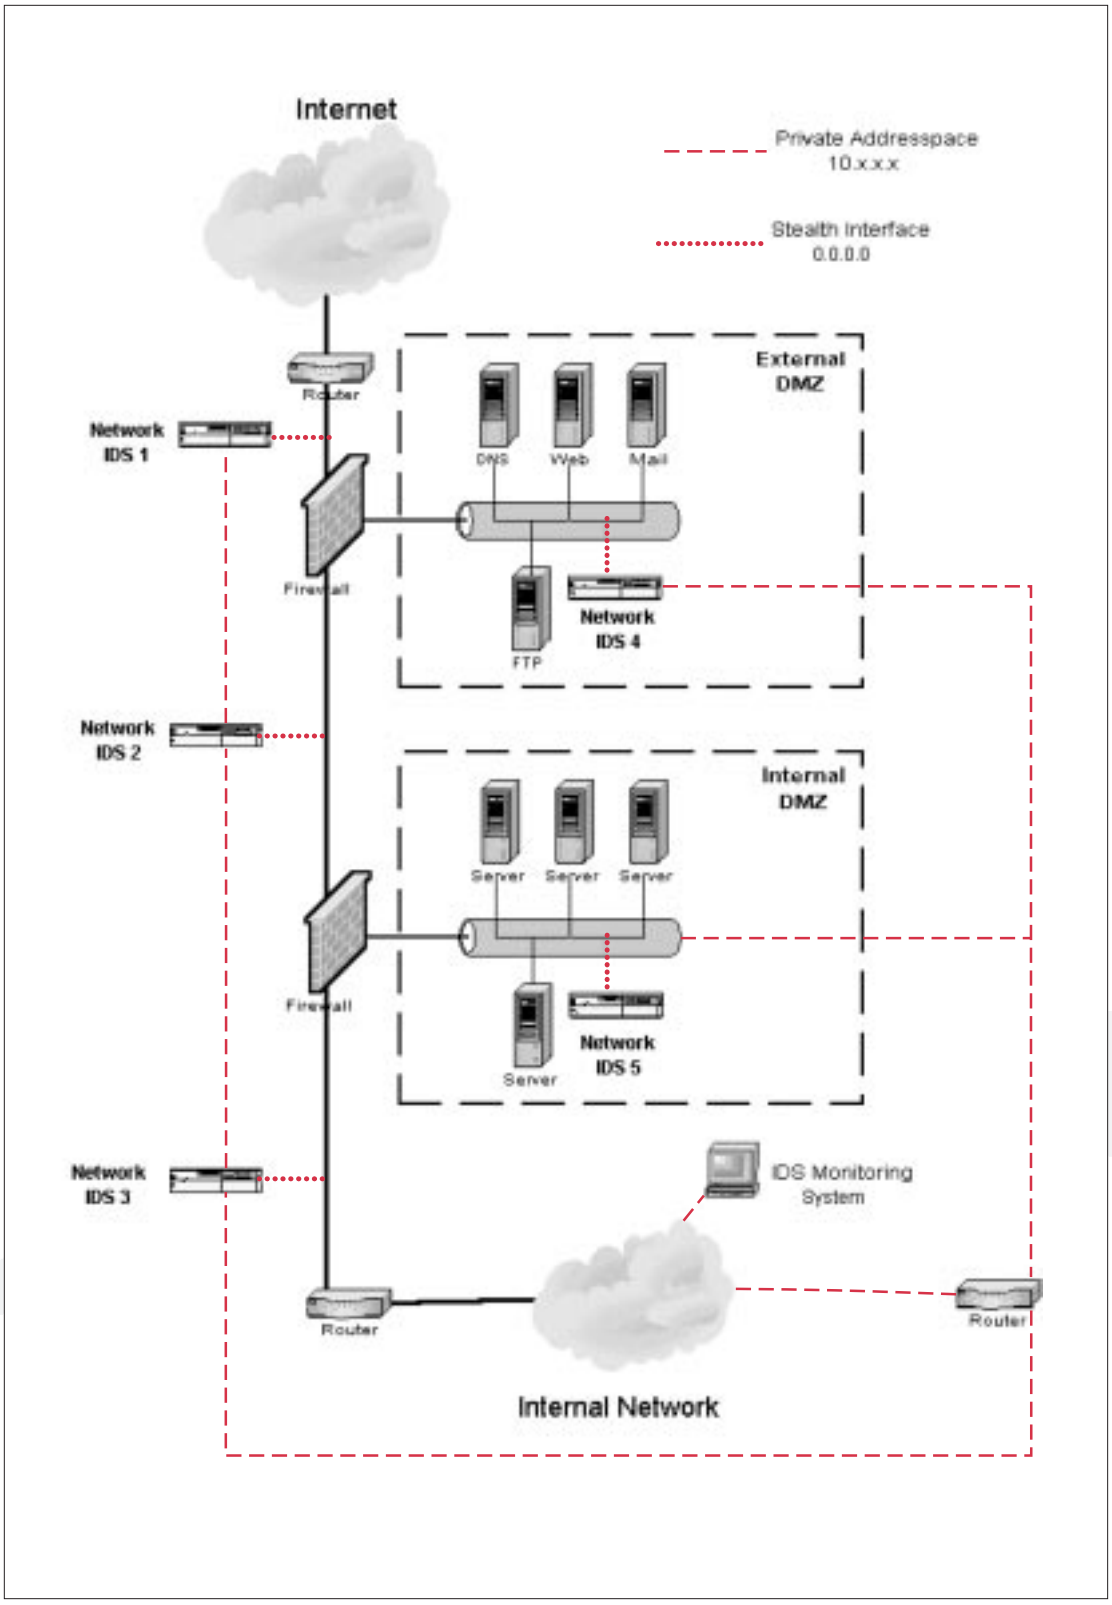

Configuring IDS Filtering to Reduce False Positives

The number of false positives generated by IDS remains the most common complaint regarding use of this technology. Two factors can affect this number and, thereby, the usefulness of IDS:

1. Placement of the sensors on the network; and
2. Tuning of IDS policies according to threat, not vulnerabilities.

Placement of the sensor on the network

As we've seen in the example above, a sensor placed outside the external firewall will generate a tremendous number of detections. So, eliminating this placement is a good strategy for reducing excess data. Only in the event of a concrete policy directive should such a sensor be used.

If, however, policy does call for IDS outside the firewall, the organization should consider these configuration rules:

1. DON'T *alert* if a network scan is observed. Rather, generate a log entry. Evidence of a scan does not, of itself, indicate an attempted break-in. Nevertheless, the system should keep track of these events for troubleshooting, investigational, or evidentiary purposes.
2. DON'T *alert* on Windows™ backdoor attacks (e.g., Back Orifice and Netbus) unless the firewall operates on a Microsoft Windows™ platform. In all other cases, the backdoor alert doesn't indicate an intrusion through the firewall. Again, log these events for later.
3. DO *alert* on DoS attacks and attacks (of any kind) directed against the firewall itself.
4. DO *log* "interesting" traffic, such as http, common gateway interface (CGI), e-mail, and remote procedure call (RPC) attacks if those services operate inside the firewall. However, DON'T *alert* on these events.

Tuning of IDS policies according to threat, not on vulnerabilities

The basic rule of IDS tuning is, "*Alert* in real-time only on events that you will act on in real-time." Judicious application of this rule will save large amounts of time and ensure the effectiveness of an IDS. For sensors behind the firewall, consider these rules:

- *Alert* on all activity indicative of Microsoft Windows backdoor types of attack. This attack signature is especially suspicious when detected behind the firewall with an external source or destination address. It requires immediate investigation.
- *Alert* on all DoS signatures, but make certain the package used does not generate false positives due to system management systems such as OpenView, CA Unicenter, or Tivoli. In particular, signatures that look for "smurf," "pingflood," and "synflood" attacks can generate false positives in conjunction with these network management systems.
- *Alert* on all buffer overflow attacks. Even if the organization does not use or plan to use the affected operating system, looking for these attacks will signal outbound traffic indicative of a "leapfrog" hacker. Observation of such traffic means that the network was compromised.
- *Log* all scanning, probing, and network mapping.
- *Log* all other "suspicious activity." Examples include traceroutes, sourceroutes, IP fragmentation, and IP duplicates.

Maintaining the IDS Technology

To keep up with the ever-changing threat environment, vendors must constantly update IDS technology. As IDS becomes a mission-critical element of an organization's security program, it must implement a process to maintain it. Otherwise, the organization will not keep up with improvements in performance, efficiency, and effectiveness.

When a particularly aggressive attack capability is discovered, for example, the difference between protecting the organization or letting it fall victim may literally depend on the number of hours it takes to install a new signature. Nevertheless, installing upgrades to IDS must follow the same standard of care as any other network modification. More and more, keeping up has placed a large burden on many businesses using IDS.

However, while upgrades are extremely important, other modifications to IDS solutions require more profound changes.

For example, the recent adoption of gigabit networking in many companies has created new challenges. Meeting these will likely require more than an upgrade to the sensors. These challenges will push IDS into previously uncharted territory. In fact, the notion of network IDS will have to give way to:

- Host-based IDS and hybrids; and
- Application switches.

These strategic approaches effectively increase the number of IDS sensors and, consequently, the number and types of possible alerts requiring monitoring. Let's look at each more closely.

Host-based IDS and Hybrids

The most modern IDS solutions integrate network and host-based IDS. Both are managed and monitored from a single system. Using a mix of both network and host-based IDS products yields the optimal IDS deployment, especially for high-speed networked environments. Two types of host-based IDS are available: traditional file integrity checking IDS and IP stack IDS.

File integrity checking type host-based IDS monitors "critical" system files and alerts when the files have been altered. This is useful, for example, in protecting web pages or user ID databases. The pioneer in this field was *tripwire*, which monitors a pre-defined list of files and ensures their integrity against a preset baseline at timed intervals. This approach has the advantage of enabling an organization to catch improper actions by insiders. These insiders can, generally, do the most damage to a system, but traditional network IDS or firewalls cannot catch them.

A recent development in host-based IDS implements IP stack sensors. These listen to all network traffic addressed to the system and compare it to a signature database, much the same way that traditional network IDS operates. The advantage of the IP stack approach is that it is only applied to the traffic targeted to the system on which it operates. One major plus of this technology is that the speed or type of networks no longer matters.

Hybrids are much easier to tune because they only listen for traffic coming to them. Moreover, the targeted approach means the sensor deals with far less traffic. This, combined with ease of tuning, makes the host-based hybrid approach much more scalable while producing fewer false positives.

Of course, care must be taken when implementing host-based technology. Some operating systems exhibit erratic behavior if their IP stack is modified. It certainly isn't an advantage to install IDS if it causes the system to crash. In the future, we can expect hardware vendors to certify host-based IDS systems, or even to incorporate them into their code base.

Application Switches

This is a fairly new way of dealing with IDS in gigabit network environments. It works by distributing traffic based on application type. The application switch keeps track of conversations enabling one to monitor accordingly. Keeping track of both sides of a conversation is especially useful for detecting more sophisticated attacks.

For example, consider a gigabit network environment hosting a web farm and multi-organization domain name system (DNS) servers. Assume a load balance of 50 percent web, 25 percent DNS, and 25 percent other types of traffic. By implementing application switching, we can dedicate one IDS sensor to DNS traffic, two sensors to web traffic, and one sensor to monitor all other traffic. Testing has demonstrated that an application switch with eight or nine IDS sensors can effectively monitor and detect all attacks, even at 90 percent utilization on a gigabit network.

Allocating Appropriate Staff

As the importance of IDS in a security program increases, an organization must develop an approach to staffing it. Designing and maintaining IDS systems requires detailed knowledge of the technology and how it works. Monitoring requires 7X24 execution of effective, repeatable processes. Incident response requires a planned, coordinated reaction on the part of many individuals.

Most organizations are finding that security is far from their core competency. Rather than going the course of trying to develop a security operations staff, many turn to external security providers. This approach has several advantages:

- Reduced personnel costs;
- Improved response capability;
- Better sensor and signature maintenance;
- More up-to-date systems; and
- Higher levels of design and configuration expertise.

The Future of IDS

Over the next three years, IDS integration of various security tools and platforms will become much tighter. The move toward host-based systems will continue, with IDS monitoring encompassing more and more of the security tasks done independently today. For example, many major virus-detection software manufacturers are working to enable remote monitoring and maintenance capabilities for the desktop. Combining IDS and virus technology would yield a big win for most organizations.

Some IDS vendors are working with firewall manufacturers to integrate these systems. By implementing the same standard of monitoring necessary for IDS to the firewall, it is possible to get even earlier warning of attacks and problems. More importantly, tighter integration of the sensing and reaction components of security (IDS and firewalls, respectively) will enable the creation of adaptive security. Firewall rules could change, automatically and in real time, to respond to an intrusion detected by an IDS sensor.

Another development in IDS technology is the adoption of 'signature-less' detection methodologies. The first working prototypes of so-called Anomaly Detection (AD) systems are in testing. These AD systems operate by learning normal network patterns within the subject environment. When run in detection mode they then alert on deviations from normal traffic patterns. While a promising concept, AD still has a number of hurdles to clear before it becomes useful for most organizations.

The Vigilinx Managed Security Services (MSS) Approach

Vigilinx has developed an MSS approach to address the issues discussed above. By assembling a team of experienced IDS designers, engineers, and administrators, we have created the most effective co-sourced solution in the marketplace. Vigilinx:

- Understands the current state of IDS technology;
- Knows how to deploy IDS sensors to the most appropriate locations within the organization;
- Identifies, configures, and tunes filtering to reduce false positives;
- Maintains the IDS technology through implementation of new features and system upgrades in real time; and
- Tracks developments in security technology, enhances service features, and develops new service products that take advantage of improved security capabilities.

Vigilinx is delivering these services today to many clients in complex environments prone to attack. It has created a state-of-the-art, proactive intrusion analysis system that provides real-time (and near real-time) event correlation, in addition to 'normal' intrusion detection analysis. This system works in real time with Vigilinx Security Intelligence Services (VSIS) to automatically reconfigure its triggers based on reports of new threats and vulnerabilities. For example, if VSIS identifies a new Eastern European hacker group targeting U.S. financial institutions, it will reconfigure monitoring and analysis systems to generate high-priority alerts for all U.S. financial customers.

To find out more about how you can make Vigilinx your organization's security team, contact your Vigilinx account representative or call us at 866.481.4101.

Conclusion

Information security is an ever-moving target. Organizations cannot underestimate the importance of looking at the big picture when implementing security systems in their environment. They must ask the question, "What makes sense for us?" Business use dictates different security strategies and required protection for systems and networks.

IDS will become an increasingly more important component of any security program, but its potential for enhancing overall security is balanced by the difficulty of its implementation and integration. Many issues face an organization just embarking on its use: design questions, information overload, performance issues, and overall effectiveness.

Each of these can be solved, but it will take expertise and time. One of the most attractive approaches to IDS is to team with a security co-sourcer such as Vigilinx. In so doing, organizations can obtain the expertise they require, effective and efficient monitoring, and the peace of mind that security is done right, enabling them to pursue their own business core competencies.

Vigilinx Digital Security Solutions

Complete strategic security services, backed by the most advanced intelligence available anywhere. Vigilinx, the clear leader in digital security solutions.

Vigilinx Security Intelligence Service™

Vigilinx Business Security Assessment™

Managed Security Services
- Intrusion Detection
- Managed Firewall
- VPN

Technology Risk Assessment

Penetration Testing

Security Architecture Design

Integration Services

Information Security

Emergency Response & Forensic Services

Proactive Forensic Services

Corporate Headquarters

45 Waterview Boulevard
Parsippany, NJ 07054

866.481.4101
973.541.5400

www.Vigilinx.com
info@vigilinx.com

Regional Offices

New York, NY

Washington DC

Los Angeles, CA

Columbus, OH

Atlanta, GA

Minneapolis, MN

Dallas, TX

Vigilinx
Digital Security Solutions