The Four Key Qualities of Effective Host Intrusion Prevention (HIP) Solutions:

Defining Deep HIP

Organizations that need to protect sensitive information assets – in order to comply with corporate or regulatory policies, protect competitive advantage, or simply enable new business processes – have come to recognize Host Intrusion Prevention (HIP) as a critical component of a defense in-depth security strategy.

This white paper explains what to look for in HIP products, and introduces the concept of “Deep HIP” as a means of characterizing effective solutions in this area.

TABLE OF CONTENTS

Executive Summary

Unrelenting and increasingly sophisticated attacks against enterprise networks have dramatically raised organizations' IT security risks. With the relative ease that many types of attacks by-pass perimeter security, traditional perimeter based security approaches are no longer sufficient to adequately protect enterprise assets. To combat these threats, security professionals are implementing multi-layered defenses, with the last line of defense being implemented at the host itself.

Host Intrusion Prevention (HIP) is the last line of defense in a comprehensive defense-in-depth security strategy. While the need for this last layer of defense is becoming increasing evident, there remains considerable confusion over what constitutes a HIP product. To be practical, HIP should be viewed as security capabilities deployed at the host to effectively keep it running, free from viruses, worms or other malware.

The key to the debate is overall effectiveness. Individually there are many HIP technologies that offer value, but do not go far enough in solving the overall problem. HIP solutions need to embody the following characteristics, or be relegated to the shelf as impractical. They must:

- Provide comprehensive protection
- Have minimal performance impact on the host
- Be extremely robust and reliable
- Offer low cost of ownership

Solutions with these attributes can offer a deep level of protection, deep within the network where an organization's most valuable information assets reside. Products with these types of capabilities provide "Deep HIP" and are critical to an effective HIP approach.

With an organization's regulatory compliance, good corporate reputation, brand equity and customer satisfaction at stake, it is imperative that organizations choose an effective solution based on the characteristics of Deep HIP as their last line of defense.

Today's Security Best Practices

Traditional network security, firewalls, and intrusion detection systems are important elements of a secure computing environment but insufficient on their own. Given the ways that the traditional network perimeter can be breached, today's security best practices implement a defense-in-depth strategy to protect organizations from attack.

> "The enterprise perimeter has changed greatly during the past several years. It now passes through mobile devices, because laptops and PDAs (personal digital assistants) often are used outside the corporate firewall. Wireless LANs allow external connections that bypass firewalls. The increasing use of Secure Sockets Layer, particularly as part of Web services, and other forms of encryption (session and data) can blind perimeter firewalls and intrusion prevention systems. Threats have changed, with rapidly propagating worms causing tremendous costs to enterprises when these worms spread across the infrastructure."
>
> **Gartner, Inc., "Management Update: It's Time for Host-Based Security Platforms"**

A Defense-in-Depth Strategy is Imperative

Maintaining a defense-in-depth strategy assumes that no single component, policy or process can assure security. The modern computing environment is too complex and diverse. Attackers exploit vulnerabilities before vendors acknowledge them or provide patches. They also have access to the same vulnerability bulletins as everyone else, and a growing range of automated tools with which to exploit them before organizations can push out a patch. The potential risk of failure and regulatory penalties require security managers not just to arm themselves against a minimum standard of documented threats but to anticipate the unknown: in effect, to 'prove a negative', and show they are not insecure.

Defense-in-depth is a dynamic process, involving a continuing cycle of risk assessment, response, and evaluation. An initial threat, risk and security audit, with special attention to servers with critical information establishes a security baseline. Once that is established, a solid defense-in-depth strategy can be created.

Economics of the Shrinking Perimeter

One common approach to defense-in-depth has been to employ the same perimeter security techniques to continually shrinking security zones. From a security perspective this is advantageous because it introduces layers of defense as well as providing the ability to tune the control to the specific needs of the asset or assets being protected. However, the economics of this approach have a meaningful impact on the nature and scope of these controls. As the perimeter shrinks the use of hardware based solutions to protect smaller and smaller zones becomes too costly and at some point necessitates a software based approach. Additionally, while the size of the zones shrinks the number of zones increases, putting an increased value on the ability to centrally manage large number of zones in a cost effective way. Taken to the extreme, the perimeter shrinks to the boundary of the host itself.

Host Intrusion Prevention is Your Best, Last Line of Defense

With the relative ease that many types of attacks by-pass perimeter security, security professionals are implementing multi-layered defenses with the last line of defense implemented at the host itself. This last line of defense is the role that Host Intrusion Prevention (HIP) solutions play in a comprehensive defense-in-depth security strategy.

> "Host-based intrusion prevention (HIP) systems have exploded onto the market to preserve the integrity of server configurations against network attacks. Network-borne intrusions have demonstrated the ability to easily pass through layers of network defenses before slamming into the hosts that are serving critical business applications over the Web…Business demands a solution that will protect application revenue streams and ensure compliance with corporate audit requirements."
>
> **The Yankee Group, Phebe Waterfield, March, 2005**

Battleground: Where Does HIP Make Sense

The purpose of information security programs and controls is to cost-effectively reduce security-related business risks to an acceptable level. When using risk as the measure to determine where HIP makes the most sense, organizations are looking at:

- High threat environments such as the DMZ where probes and attacks are frequent
- High value hosts where the value of the asset either to the organization or the attacker is significant and warrants additional controls
- Hard to patch environments that remain vulnerable longer because they are not easily patched

Confusion Surrounding HIP

While the need to provide a last layer of defense at the host itself is an easily understood problem, there has been considerable confusion over what constitutes a HIP product. Security vendors and analysts have all jumped into the fray, each positioning a slightly different view of what constitutes HIP technology, including existing technology such as firewalls, Intrusion Detection Systems (IDS) and anti-virus signature based approaches (Figure 1).

Figure 1: Competing & Overlapping Definitions of HIP

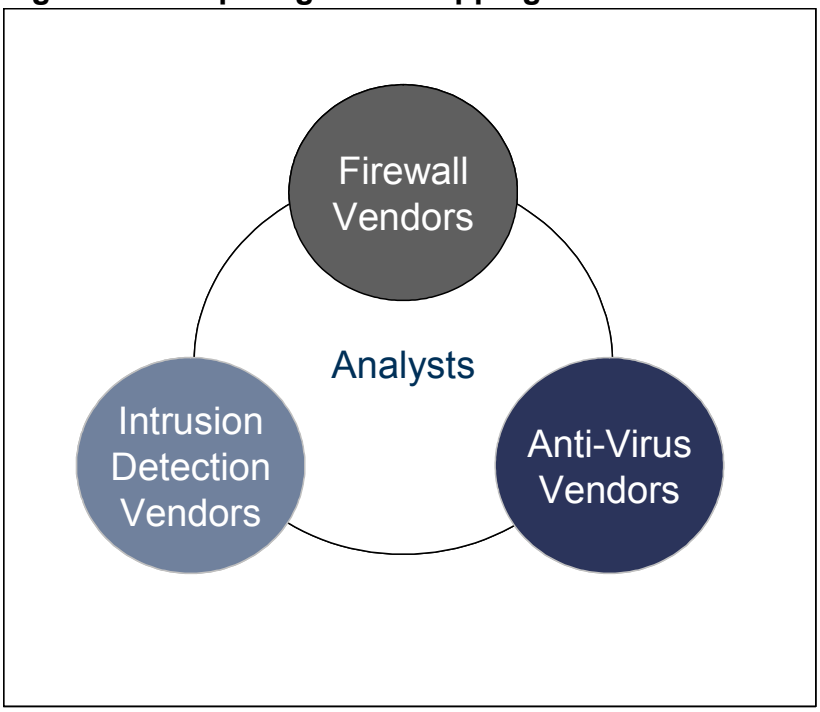

Even among analysts there are varying definitions of HIP. According to Gartner VP and Distinguished Analyst, John Pescatore, “HIP systems detect and block malicious operations and attacks, and do so without disrupting normal operations. Hosts running intrusion prevention processes will only show latencies in the tens of milliseconds range, even when attacks are being blocked. To be effective, HIP systems must use a number of algorithms, based on rules, policies and anomaly detection as well as signatures, operating at the application as well as network levels. A true HIP system does not simply restrict the set of applications that can execute on a server; it also blocks anomalous behavior in the applications that are allowed.”

The Yankee Group use a slightly different definition, identifying the following key attributes of a HIP solution:

- It executes after signature-based filters have examined data
- It finds deviant behavior, making ‘block and allow’ decisions based on knowledge of normal application behavior
- It ‘learns’ the range of acceptable application behavior over time
- It blocks abnormal behavior in real time
- It delays attack propagation by timely incident notification through a management console

Unfortunately these overlapping, competing definitions create confusion for organizations. To be practical, HIP should be considered in the context of the problem

that organizations are trying to solve. Put simply, HIP should be viewed as security capabilities deployed at the host to effectively keep it running, free from viruses, worms or other malware.

The Four Key Qualities of Effective HIP

This broad definition provides a lot of leeway in what would be considered HIP technology. The key, however, is overall effectiveness. HIP solutions need to embody the following characteristics, or be relegated to the shelf as impractical:

- Comprehensive Protection
- High Performance
- Robust Security
- Low Cost of Ownership

1. Comprehensive Protection

Effective HIP solutions need to offer comprehensive protection in order for them to be of value in today's computing environments. Many first generation approaches suffered from being too specific in terms of the platforms they protect or the types of threats and vulnerabilities they protect against.

Breadth of Host Coverage

Given the heterogeneous nature of computing environments and the range of threats and vulnerabilities, HIP solutions need to provide breadth of host coverage in two very distinct ways. They need to cover a broad range of host platforms as well as being able to cover a wide range of target applications.

Breadth of Attack Coverage

With new vulnerabilities and corresponding exploits being created all the time, breadth of attack coverage is critical not just for providing adequate protection for today's attacks, but also to protect against future attacks.

A key consideration that organizations should take into account when looking at attack coverage is level of risk. Any two attacks or classes of attack are not necessarily equal in terms of the risk they pose to a host, and should not be considered equal when evaluating HIP products. For instance, many types of attack need a certain sequence of events to occur for the attack to be successful. In many host environments, the specified sequence of events for one type of attack may already be prevented by other protection mechanisms or the probability of them occurring may be extremely low when compared with other types of attack.

Remote versus local attacks are a classic example of this challenge. Local attacks, even if the malware was transferred from a remote source, requires a specific action to occur at the host before the attack is successful, such as running a specific service or

surfing to a specific web site. In many cases, organizations have existing protective measures in place for their servers, such as physical security, operator training or configuration management that already greatly reduce or eliminate the risk posed by these types of attacks. Most remote attacks, on the other hand, are not prevented by such measures and pose a much more significant risk to server environments. That is why effective HIP products should pay particularly close attention to the range of remote attacks.

There are a number of organizations providing valuable insight into classifying attacks and providing information to help enterprises assess their risk. In particular, enterprises should consider classifications provided by the Open Web Application Security Project (www.owasp.org) and the Web Application Security Consortium (www.webappsec.org) as a starting point for evaluating product coverage.

Accuracy

No security control can be considered comprehensive if it is not accurate. Accuracy is usually measured based on ability to prevent false negatives and false positives. False negatives in the context of HIP are instances where malicious system execution or data traffic is not prevented by the control and allowed to occur. False positives on the other hand, are instances where the control prevents legitimate system execution or data traffic.

For many types of security controls the two measures of accuracy are inversely related. As the sensitivity of a control is increased to lower the incidence of false negatives (or what is often referred to as the *false acceptance rate*), the incidence of false positives (or *false rejection rate*) increases. Conversely when control sensitivity is lowered in order to reduce the incidence of false positives, the incidence of false negatives increases.

Both from security and operational perspectives, neither false positives nor false negatives are desirable. While the security group may be naturally driven to dialing down false negatives, the increased number of false positives results in more incidents to investigate and operational issues as more and more legitimate transactions are blocked. Similarly, driving down false positives increases the probability of a breach and any consequences thereof. In the end, both false positives and false negatives drive overall operational cost. Effective HIP solutions (Figure 2 – “HIP Solution 2”) should delivery superior accuracy in terms of zero, or near-zero, incidents of false positives and false negatives. As well, an effective HIP solution provides a broader operating range, which reduces the need for continually tuning the system.

Figure 2: The Advantages of Improved Accuracy

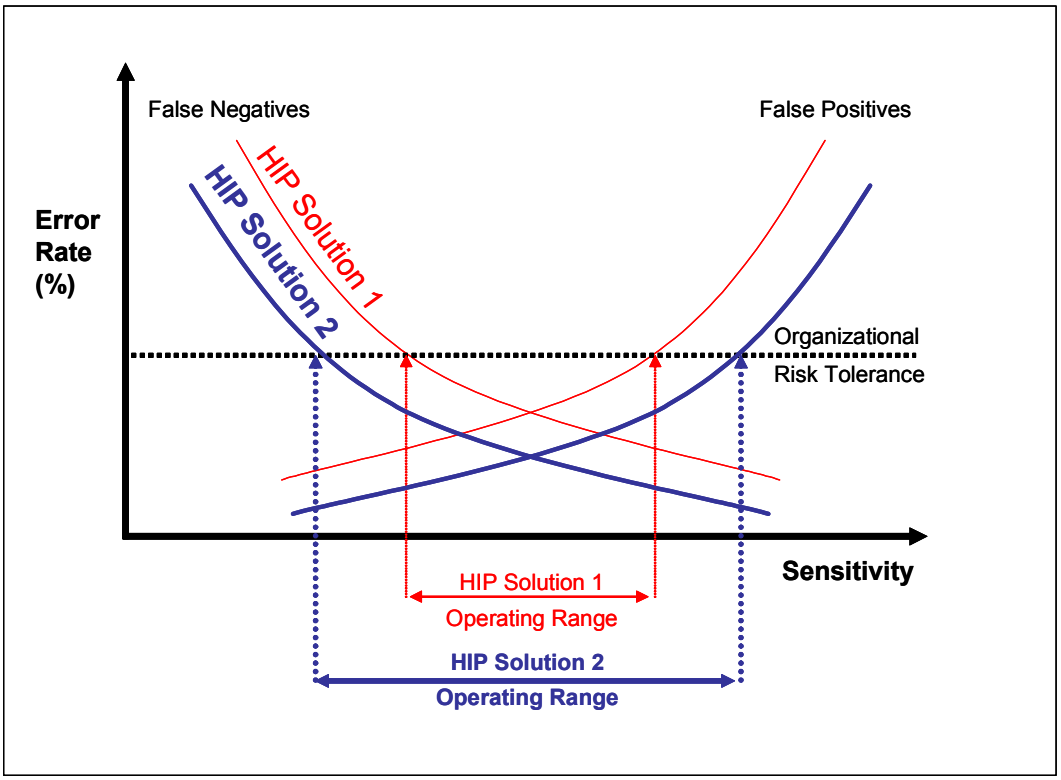

2. High Performance

While it is important to provide excellent protection at the host, if it comes at the cost of host performance it will either not be used or significantly add to the overall cost as system architecture is modified to compensate.

To avoid this problem, HIP solutions should not just provide high throughput but consume only a small fraction of host resources and behave in a predictable way – allowing the operations group the confidence that the host will continue to operate as desired while at the same time ensuring they operate securely.

3. Robust Security

Given that the attackers' objective is to compromise the host, one of their first tasks is to either evade or disable any security mechanisms in place preventing compromise or otherwise limiting the value of the target to the attacker. In order to be useful, HIP solution needs to be robust and provide proactive protection.

An effective HIP solution should eliminate threats and defeat attacks before they have had the chance to penetrate the host. Specifically, effective solutions should provide in-line protection as close to the network layer as possible and eliminate both known and unknown attacks. Many approaches such as signature based anti-virus protection only protect against known attacks where a signature has been published. Attacks for which signatures have not yet been developed, but that possibly exploit the same weakness are not protected against.

Another way that Deep HIP products can provide proactive protection is to be implemented in a way that insolates itself from direct attack.

4. Low Cost of Ownership

With the overall goal of a security control being to cost-effectively reduce security-related business risks to an acceptable level, organizations need to concern themselves not only with the risk-related costs that are being mitigated, but also with the cost of the control itself. HIP solutions need to offer low cost of ownership in order to make sense in the overall risk equation.

In addition to factoring in the acquisition and maintenance costs of the solutions themselves, organizations should be looking at the overall operating impact. In many cases, it has been the operating impact of the solution that has negatively impacted the perception of many HIP solutions. Several operating characteristics that can have a high impact on cost include:

- Accuracy – if a solution doesn't provide the needed accuracy, responding to false positives or negatives will drive up operating costs.
- Performance – if the performance impact is significant it may require additional expenditures in order for the overall system to operate as needed.
- Operational Intrusiveness – if a solution directly impacts "normal operations" it can significantly drive up direct and indirect costs of the solution - ideally a solution would minimize operational impacts in order to reduce its cost impact.

Operational Impact

HIP products introduce another variable into the already complex host equation - one that needs to work well with all the other solution components and the operational processes in place. To be non-intrusive, effective HIP products should:

- Be insulated from dynamic components of the system such as the application and OS so as not to introduce complexity or limit upgrade paths.
- Be easily testable and provide the ability to quickly rollback to a previously known state to facilitate system troubleshooting and meet change management requirements.
- Provide central management capabilities that can easily scale to large number of hosts.

- Be easily tunable as the threat environment changes and offer a relatively flat acceptable operating range

Conversely, effective HIP solutions should not:

- Take weeks or months to train on normal system behavior or need retraining whenever a change on the host occurs in order to be effective
- Require service or machine reboot in order to take effect

Learning from First Generation HIP Approaches

There have been many products and technologies that have been positioned as HIP over the last few years. In many respects these represented the first generation of HIP solutions and can be broadly classified into two general approaches; System Execution Control and Application Proxy Data Filtering.

Application Proxy Data Filtering

As the name suggests, application proxy data filtering approaches are implemented as a proxy for the application and utilize data signature approaches at the application protocol level as a means of detecting malicious data. Depending on whether the particular product was implemented as an in-line/active or a passive/reactive control, prevention is enforced by either dropping the specific request or by sending TCP resets to either end of the communication.

System Execution Control

Unlike a data inspection based approach, system execution control approaches typically use heuristics and artificial intelligence techniques to compare system behavior against a baseline that has been learned over time. Detected behavior outside of the “normal” baseline is considered malicious and prevented by controlling system calls from the application to the OS.

Analysis of First Generation Approaches

Like any first generation approach, much can be learned by looking at their limitations. Although individual products may differ, the following table outlines the general limitations of solutions based solely on these approaches.

Table 1: Limitations of First Generation HIP Approaches

	Application Proxy Data Filtering	System Execution Control
Comprehensive Protection	• Limited by the number of protocols supported • Typically only inbound control	• Doesn't communicate what it has prevented • Can't protect against exploits utilizing normal system behavior (e.g. SQL injection) • Poor accuracy • Prone to high false positives particularly in complex environments
Performance	• Implemented in user mode with full protocol decoding significantly impairing performance	• System behavior monitoring consumes considerable host resources
Robust Security	• Does not necessarily operate in-line • Relies on protocol decoding that may render it subject to the same vulnerabilities as the target system	• Control acts after attack begins execution
Low Cost of Ownership	• Significant performance impact • Not easily testable • Often requires system or service reboot	• Significant performance impact • Implemented between the OS and application • Requires significant training/retraining before effective • Often requires system or service reboot

Bringing It All Together: Defining Deep HIP

Much has been learned from first generation HIP approaches. Additionally there are a number of traditional security technologies, such as firewalls and signature based systems, when deployed at host offer HIP. Individually, many of these technologies offer value but do not go far enough in solving the overall problem. Taken together however, and designed to meet the key qualities of effective HIP, they can offer a deep level of protection, deep within your network where your most valuable information assets reside. Products offering these types of capabilities – comprehensive protection, high performance, robust security and low cost of ownership – provide "Deep HIP", and are critical to an effective HIP approach.

Protecting Your Organization: The Need to Act Now

Today, attackers can exploit a vulnerability so quickly that traditional protection is inadequate. With no patch to plug the hole, or no signature to identify and block the malware, the enemy is within the gates before you know it. And the problem is becoming more critical. Even while security managers struggle to protect their networks, senior management is demanding greater openness, through mobile computing, wireless networking, Web applications and closer online relationships with suppliers and customers.

> For improving security-related technology controls, best-in-class firms focus on segmenting access to corporate and customer data and favor brand name and best-of-breed security solutions that are added to the environment, instead of using security capabilities that come prepackaged with systems and network platforms.
>
> **AberdeenGroup, "Most Important Security Action: Limiting Access to Corporate and Customer Data", Jim Hurley, VP Research, March, 2005**

In response to these challenges, more and more organizations are turning towards Host Intrusion Prevention as a critical part of their overall information security strategy. With an organization's regulatory compliance, good corporate reputation, brand equity and customer satisfaction at stake, it is imperative that organizations choose an effective solution based on the characteristics of Deep HIP.

For organizations looking for a layered defense-in-depth strategy that includes Deep HIP, find out more about products and solutions offered by **Third Brigade,** please visit www.thirdbrigade.com.

About Third Brigade

Third Brigade specializes in providing intrusion prevention systems (IPS) to health care, government, telecommunications, financial services and other organizations that need to prevent attacks that exploit vulnerabilities in commercial and custom software, including web applications. It enables you to create and enforce comprehensive security policies that proactively protect critical applications, sensitive data, and hosts, ensure regulatory compliance, and maximize the performance of your people, processes and hosts. Unlike other intrusion prevention systems, Third Brigade's is not intrusive. It has been architected from the ground-up for intrusion prevention, and is smaller, faster and simpler.
For more information, please visit www.thirdbrigade.com, or contact us at:

Corporate Headquarters
40 Hines Rd
Suite 200
Ottawa, Ontario, Canada
K2K 2M5
Toll free: +1.866.684.7332
Local: +1.613.599.4505
Fax: +1.613.599.8191

United States Headquarters
11710 Plaza America Drive
Suite 2000
Reston, Virginia 20190
USA
Toll free: +1.866.684.7332
Local: +1.703.903.4479
Fax: +1.613.599.8191

Author Profiles

Brian O'Higgins, Chief Technology Officer

Brian is a veteran security professional, and is best known for his role in introducing PKI (Public Key Infrastructure) technology and products to the security landscape. Prior to joining Third Brigade, Brian was the co-Founder and Chief Technology Officer of Entrust, a leading Internet Security company. While at Entrust he had overall responsibility for the technology vision and direction for the company. He was previously with Nortel where he established the Secure Networks group in 1993, and was instrumental in spinning-out this group as an independent company, Entrust.

Blake Sutherland, Director Product Management

Blake has spent much of his career understanding customer and market requirements and incorporating them into software products. Prior to joining Third Brigade, Blake was at Entrust, a leading Internet Security company, as a Senior Product Manager and Solution Manager for Entrust's Secure Data Solutions. Blake is a Professional Engineer in the Province of Ontario and holds a Bachelor of Applied Science degree in Engineering Physics from Queen's University.