Intrusion Defense Firewall

Network-Level HIPS at the Endpoint

I.   EXECUTIVE SUMMARY

Mobile computers that connect directly to the Internet outside of a company's firewall can introduce risk into the corporate network and thus require a higher level of security to protect against network intrusions. Host intrusion defense systems combine intrusion detection and prevention capabilities, and run on the host itself.

There are two main approaches to host intrusion defense: 1) system execution control; or 2) a network approach. The network approach offers several advantages by blocking malicious code before it impacts the host, targeting potential vulnerabilities and known exploits, in addition to providing proactive vulnerability-facing network inspection.

Using multiple techniques to filter both inbound and outbound traffic insures optimal efficiency and effectiveness. This blended approach includes deep packet inspection, exploit and vulnerability filters, and custom filters to protect custom applications. A tunable, flexible solution insures delivery of business-critical communications while protecting against unwanted network traffic.

This white paper examines a blended approach to host intrusion defense in search optimal security combined with an ideal balance between flexibility, control, and ease of management. Specific examples highlight the benefits and pitfalls of the many different filtering techniques.

II.   THE TWO APPROACHES TO HOST INTRUSION DEFENSE

Host intrusion defense systems combine intrusion detection and prevention capabilities, and run on the host itself. They complement existing network security mechanisms, acting as another layer of protection against attacks that now routinely bypass or penetrate perimeter defenses, and target vulnerabilities in software on the host.

Although there are many different host intrusion defense systems available to enterprises, there are two main styles or approaches. These two approaches are fundamentally different but share the same objective of keeping malware off the host. The challenge for any system is to achieve a high degree of accuracy by minimizing the number of false positives (blocking good data) and false negatives (allowing bad data). There is a certain amount of tuning that is necessary to make sure the system is operating optimally.

Style 1: System Execution Control

System execution control is often referred to as a behavior-based approach. These systems learn what the "normal behavior" is for a host, and then they can identify and block strange or anomalous behavior. Typically, this approach uses techniques such as system call interception, which monitors the interaction between application software and the operating system. Most first-generation host intrusion detection and prevention systems (IDS/IPS) tend to use the system execution control approach.

The advantage of this approach is that it provides a broad protection umbrella that covers any operational anomaly. These systems can, for example, protect more than just the network interface and can cover attacks launched from portable storage devices and the keyboard. System execution control systems, by design, also do not need signature updates since they provide zero-day protection once they are trained on "normal" behavior.

The disadvantage of system execution control systems is that they have relatively high care and feeding requirements. Each host must be trained to establish the rule set and continuously be retrained as software (including operating systems and enterprise and web applications) is updated. Another maintenance issue is the removal of malicious code. Even though malicious code might have been blocked from executing, the infected machine still needs to be cleaned.

Style 2: Network Approach

A different approach to solving the same problem is the data network style. This approach uses traditional, proven network perimeter defenses such as firewall, IDS and IPS, but applies them at the network layer on the host. The enforcement point is typically kernel mode based. Although this approach has a smaller coverage umbrella compared to system execution control, it does cover the network interface, which is the attack vector of greatest concern, especially with today's increase in blended web threats. In many cases, especially with mobile laptops, it is the highest priority concern.

In contrast to system execution control, the network approach is also more proactive: it stops malicious code before it gets on the host. It can, however, be challenging to understand the packet stream in enough detail to make accurate decisions on whether the data should be allowed or blocked. Instead of training, these systems are tuned with rule updates to control the blocking. These rules and signatures are different than malware signatures used by Trend Micro's virus and spyware scan engines. They are proactive by covering the vulnerability, rather than individual exploits.

III. ADVANTAGES OF THE NETWORK APPROACH

There are many trade-offs in determining the defense approach that is not only appropriate for your environment, but also for the risks and threats you face. The two main business drivers for host intrusion defense are, however, quite universal: 1) to protect particularly exposed endpoints; and 2) to provide protection until you can patch vulnerabilities. This is becoming more and more urgent in today's world of fast moving attacks that hit before patches, assuming they're available, can be downloaded, tested, and deployed. The bulk of malicious code and targeted attacks now occur soon after a software vulnerability becomes known. The vendor's announcement of the patch update itself may start the race. You have to shield or patch as soon as possible. The vulnerability-shielding aspect of host intrusion defense offers immediate value since the shields can be updated without a system re-boot or the extensive testing required by a bundled patch update.

In short, a network approach to host intrusion defense with vulnerability-facing signature updates is an effective, proactive solution. Just having this host intrusion defense agent on exposed endpoints or on endpoints regularly handling compliance-relevant data can make your audit compliance negotiations much simpler. Logs that show blocking of specific attacks on critical and vulnerable applications are a fundamental part of demonstrating that your operations are secure, and they help justify the investment in host intrusion defense.

IV.   TREND MICRO'S BLENDED APPROACH TO HOST INTRUSION DEFENSE

The Trend Micro™ Intrusion Defense Firewall plug-in for OfficeScan™ Client/Server Edition 8.0 is an advanced host intrusion defense system that uses multiple techniques to filter malware from the incoming and outgoing traffic stream (figure 1). It is the blending of multiple filters that offers an extremely efficient and broad range of protection against malware. The layered approach can be compared to sifting gravel through a series of increasingly finer grained screens. You don't start with the fine screen because it would immediately get clogged up with larger stones.
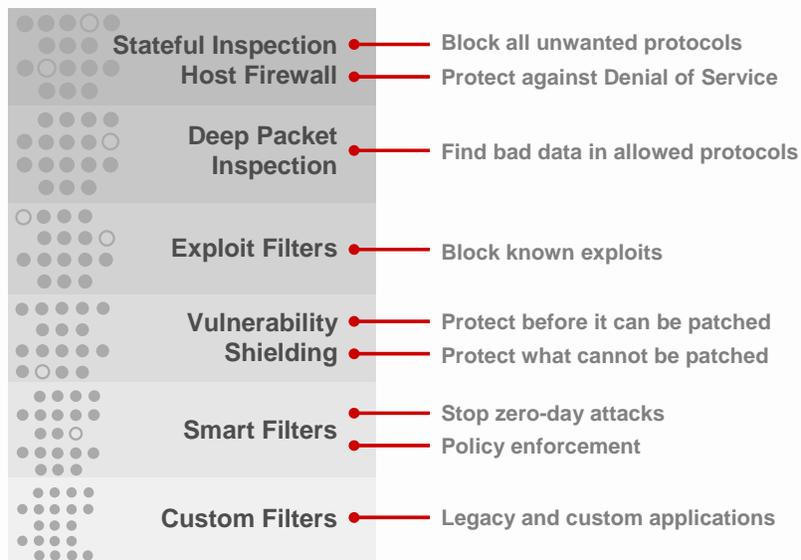
Stateful Inspection	Block all unwanted protocols
Host Firewall	Protect against Denial of Service
Deep Packet Inspection	Find bad data in allowed protocols
Exploit Filters	Block known exploits
Vulnerability Shielding	Protect before it can be patched
	Protect what cannot be patched
Smart Filters	Stop zero-day attacks
	Policy enforcement
Custom Filters	Legacy and custom applications

Figure 1: Trend Micro's Blended Approach

Step 1—Stateful Firewall

A stateful inspection packet filtering firewall allows traffic that is known to be good, and blocks everything else. This step dramatically reduces the attack surface area, as all ports are closed by default, and the firewall rules open up the specific ports required by the applications on the host.

Step 2—Deep Packet Inspection

Next, the traffic that goes through the firewall is examined with deep packet inspection technology that looks for patterns in the payload. Each byte of the packet is examined just once to minimize performance impact, and the sequence of rule sets that control deep packet inspection filtering follow in parallel. While these steps execute in parallel, they follow the logical order shown in figure 1.

Step 3—Exploit Filters

Here, known malware is efficiently detected and filtered out with exploit filters that use signatures for individual exploits that are well-known and widespread. This is similar to antivirus signature updates. A long list of specific malware signatures is not required, but selected high runner exploits get their own filter. In addition to efficiency in detection, this allows very specific reporting in the logs since the originating IP address and specific exploit can be recorded.

Step 4—Vulnerability Filters

Vulnerability-facing filters have the greatest business benefit, as one filter will shield a particular vulnerability from an unlimited number of exploits. It may turn out that a new exploit can evade a current vulnerability signature, but if that ever happens, an updated exploit filter or vulnerability filter can be deployed. The update mechanism allows revised or new filters to be pushed out automatically whenever necessary.

Step 5—Smart Filters

Smart filters provide enterprises with the ability to enforce corporate network policies for the use of certain applications. For example, administrators can control whether Instant Messaging applications are allowed, and if so, designate which Instant Messaging clients are supported. Administrators can also use smart filters to block peer-to-peer applications such as Skype and BitTorrent, and media streaming applications such as YouTube. In addition, smart filters can help determine which browsers—such as Internet Explorer, Safari, Firefox, and Opera—are supported in the enterprise. Mitigating actions include dropping the connection or selectively blocking or even modifying offending bytes in the packet.

Step 6—Custom Filters

Custom filters can be developed to provide additional protection for specific protocols, and custom and legacy applications. They can also be designed to log application security events. Unlike behavior-based systems, which often have a closed design that does not allow for customization, the Intrusion Defense Firewall's open design allows third parties and customers to create their own custom filters.

V. BENEFITS OF INTRUSION DEFENSE FIREWALL'S BLENDED APPROACH

All intrusion defense systems need to be tuned for optimal operation in order to reduce false negatives and false positives. The firewall blocks a lot of traffic, but opening the door for port 80, for example, results in false negatives as malware embedded in HTTP traffic now gets a free ride into the network. Custom filters that are designed to lock down one particular application for example, can result in false positives. As the controls get tightened to only allow specifically formatted data, there would be little chance of missing exploit code, but a greater chance of rejecting good data. Exploit filters that stop a particular exploit specimen and vulnerability filters that shield a known vulnerability offer a good balance between the two error types.

The error trade offs between a false positive and false negative also track the order in figure 1. As we move from Step 1 to Step 6, the chances of a false positive increase. As we move up from Step 6 to Step 1, the chances of a false negative increase. This again illustrates the importance of, and flexibility in, tuning. Using the right mix of filters is the secret to finding the optimal balance point.

The Intrusion Defense Firewall is bi-directional; it allows different rules to be applied to data entering or leaving the endpoint. This allows you to deal with both incoming attacks and outbound compliance issues. For example, in an e-health patient record application, a custom rule could block a specific message type containing personal information from leaving the endpoint if it is not encrypted. In normal operation these messages should be encrypted, but maybe a configuration problem on a backend system allowed this data to go out unencrypted anywhere on the Internet, instead of being encrypted for only a select set of trusted endpoints.

The Intrusion Defense Firewall plug-in for OfficeScan is unique in that it not only can allow or block data, it can also modify data. Data modification rules are used sparingly, but they can be quite effective in neutralizing potentially malicious code without taking down the session and creating a false positive. One example of a simple data modification rule is to alter the response to banner scans. This can deflect some automated attacks that are looking for the signature of a particular system. The operation and flexibility of the different filter types is explained in more detail, with examples, in the next section.

TREND MICRO

VI. FILTERING STEPS IN ACTION

Firewall

There are two important considerations when implementing a host-based firewall: 1) having comprehensive controls over inbound and outbound traffic; and 2) making it manageable.

By controlling which traffic is allowed to access and leave a host, the attack surface of the host is minimized. Implementing this with a relatively small rule set is essential to reducing management overhead and the chance of configuration errors. The Intrusion Defense Firewall rules employ an object re-use paradigm for rule construction, which allows the rule set to be compact. Restrictions on source and destination MAC and IP addresses can be used to ensure traffic is only coming from trusted hosts.

An example to consider for rule changes is the following: imagine you have two separate network segments (A and B) each containing 100 endpoints. Firewall restrictions in each zone allow each host to talk to hosts within its segment while outside connections are restricted to a limited number of hosts. Now you want to move five endpoints from segment A to segment B. In many systems this would represent up to 10 rule changes at up to 200 different hosts. However, with Trend Micro's centrally managed system, this can be accomplished with as little as two rule changes, along with a new "segment B policy" change to the five servers.

Firewalls are a commodity today, and many host firewalls that are included with the operating system, such as IPTables, provide excellent protection. However, they do not provide centralized management, which is required to make them cost-effective to operate and maintain. When you consider the scale that is implied with host-based firewalls, ease of management is essential.

Exploit Filters

Exploit filters provide protection based on the characteristics of a known exploit against a known vulnerability. In some cases, an exploit may be unique in its method of attack on a vulnerability. In this case, exploit filters are simply the most efficient way to detect and block it. In addition, exploit filters can support reporting and audit requirements by providing information on how many times a particular attack has been launched against the organization.

Vulnerability Filters

Vulnerability filters shield known vulnerabilities from unknown exploits. This is the filtering step that currently provides the maximum business benefit for the Intrusion Defense Firewall. An example of a vulnerability that is best addressed by this type of filter is the Windows™ Metafile (WMF) vulnerability in the Windows Graphics Rendering Engine, which allows for arbitrary code execution (MS06-001).
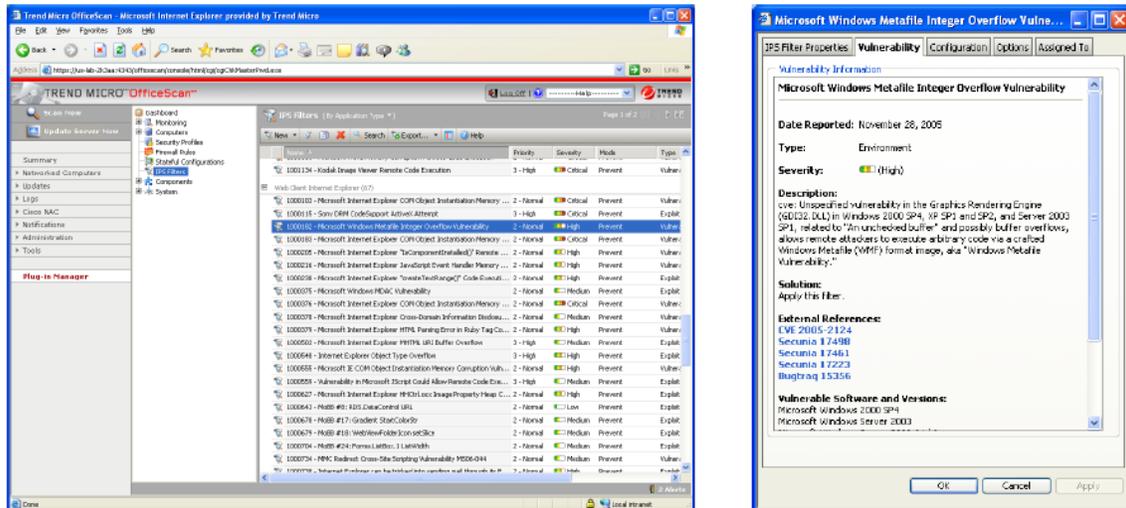
Figure 2: Windows Metafile Integer Overflow Vulnerability Filter

A Windows Metafile is a standard Windows image file format. It consists of a set of graphics functions and parameters that describe the steps for rendering an image. A WMF has a 16-bit format that can contain both vector and bitmap information. A WMF file contains a header followed by one or more records of data. Each record is a binary-encoded function call to the Microsoft Windows Graphics Device Interface (GDI). The data from each record is passed to the respective GDI functions as a parameter to render the desired image.

One such function is the SETABOPRTPROC function, which sets the application-defined abort function that allows a print job to be cancelled during spooling. The second argument expected by the function is a pointer to an arbitrary function. When a WMF file calls this function, the function code is directly supplied as the last parameter. The first parameter is skipped due to the defined calling convention for WMF format.

The vulnerability in this case exists in the Microsoft™ Windows™ operating system core graphics component. The vulnerability can be exploited by the unlimited access to GDI functions provided to WMF files. Specifically, the ability to invoke the SETABORTPROC GDI function allows a WMF file to deliver arbitrary code that is called by the operating system.

To exploit this vulnerability, an attacker may deliver a malicious WMF file, and once the target user is persuaded to open the malicious resource, the vulnerability is triggered. The code (malicious payload) delivered is used when the abort procedure is called either because an error occurs during the processing of the WMF file or as a result of an explicit GDI call.

Unlike an exploit filter approach that would focus on identifying known malicious payloads, a vulnerability filter approach looks for specific start and end patterns that indicate the presence of a Windows Metafile and then checks for calls to SETRBABORTPROC or MFCOMMENT, which would indicate that the vulnerability is being targeted.

Smart Filters

Smart filters are designed to provide network-based application control to help enforce corporate network policies. Smart filters can be deployed in either detect mode or in prevent mode (for blocking). For example, the Application Control for Opera Web Browser filter could be configured to generate an alert whenever Opera browser traffic is detected on the network, but to allow the user to continue to use the Opera browser. These alert intervals can be configured to an optimal timeframe, such as alerting once a day. Alternatively, the Application Control for YouTube filter could be configured to block any media streaming requests from the YouTube website so that users cannot view video clips when at work.

Custom Filters

Custom filters are smart filters that are tailored for a specific application, rather than distributed as a generic package. Because the Intrusion Defense Firewall plug-in provides an open design, these filters may be created by customers. In many cases they are simply modifications of existing filters that accommodate something unique for an application.

VII. CONCLUSION

Any host intrusion defense system requires precise balancing between false positives and false negatives for optimal operation. Trend Micro's approach is unique because of the multiple filtering techniques applied at the network layer.

Intrusion Defense Firewall Plug-In for OfficeScan enhances endpoint security with a Host Intrusion Prevention System (HIPS) designed to protect applications and systems from vulnerability exploits, and shield vulnerabilities from threats before patches can be deployed. The extensible plug-in architecture literally extends the OfficeScan lifecycle, requiring less administrative effort than full product updates. As a simple plug-in, it is easily deployed and managed within the existing infrastructure. The solution's high-performance, deep-packet inspection engine monitors incoming and outgoing traffic for network protocol deviations, suspicious content that signals an attack, or security policy violations. The result is an extremely effective solution consisting of a firewall and five types of complementary filters with tuning flexibility.