Realtime publishers

"Leading the Conversation"

The Essentials Series: Selecting the Right Network Threat Management Solution

Intrusion Prevention Evaluations and the Perils of Checkbox Product Comparisons

sponsored by

NOKIA

by Dan Sullivan

Intrusion Prevention Evaluations and the Perils of Checkbox Product Comparisons.1
The Past: Too Many Promises, Too Little Delivery....2
Excessive Logging....2
Too Many False Alarms....3
Difficulties Scaling....3
Adverse Impacts on Business Operations....3
More Effective IPS Evaluations....4
More Effective Intrusion Prevention Deployments....4
Summary....5

Copyright Statement

Intrusion Prevention Evaluations and the Perils of Checkbox Product Comparisons

Traditional intrusion prevention systems (IPSs) have been championed by advocates, while critics have found faults and unmet expectations. An obvious question is, who are we to believe? Actually, a better question is, how are we to evaluate IPSs? This question in turn is driven by several business and technical factors:

- What are the business and risk management drivers behind the purchase of an IPS?
- What kind of information is generated by the IPS once it is installed and operational?
- What is required to tune and optimize the performance of an IPS in your environment?
- How can information gathered by an IPS be leveraged to assist with network management?

IPSs serve a specific purpose and have particular characteristics not shared with other security controls. Evaluations of IPS technology should be adapted to these characteristics. This article will examine:

- Unmet promises of IPSs
- The adverse impact on business of these unmet expectations
- A broader method for more effective IPS evaluations

The goal of this article is to identify the reasons for past shortcomings and to understand how to effectively evaluate IPS technology so that decision makers can be confident that their expectations will be met.

The Past: Too Many Promises, Too Little Delivery

Early adopters of intrusion prevention technology were promised many things: the ability to detect anomalous activity on the network, shut down rogue connections, and protect the network from attacks directed at critical servers. Things did not always work out as promised. It is not that there were fundamental flaws in the technology, but implementing IPS operations with the surgical precision needed to eliminate real threats without disrupting normal business operations proved more difficult in practice than was apparent at first.

Broadly speaking, three types of problems occurred too frequently in early IPS implementations:

- Excessive logging
- Too many false alarms
- Difficulties with scaling

Taken together, these drawbacks left IPS implementers with less than apparently promised.

Excessive Logging

IPS applications can collect large volumes of logging detail. This is understandable. There are many protocols used on a typical network to communicate between many varied devices about a wide array of services. Many of the patterns in this traffic match patterns in attack signature databases, leading to alerts. Unfortunately, too much raw log data and insufficient filtering and reporting can cause an IPS to actually fail to address the problem of information overload facing network administrators. IPS data that is too low level can be a distraction because it is not linked to decisions or actions that systems managers can make to improve the state of their network security. Ideally, the information provided by the IPS should take into account the context of events and provide information about relevant events with appropriate levels of detail.

Too Many False Alarms

IPSs have been criticized for another kind of excess: too many false alarms. Alarms are generated when an anomalous event occurs on the network. Ideally, anomalous events are triggered when a potentially problematic event occurs and such an event is consistent with what is known about the state of the network. For example, if a large volume of data is transmitted from a database server to an external IP address at a time of day when there is usually little database activity and the event details are consistent—for example, the database server is located on that segment—then the alarm should be reported. Such an alarm fits with the network context and expected events. Similarly, context can be used to filter false alarms. If an alarm on a Microsoft Exchange Server is generated from a network segment that does not have an Exchange Server, it can be safely filtered as a false alarm.

Difficulties Scaling

Another potential problem with IPSs is scaling to meet the demands of the network. IPSs use large rule bases and statistical pattern recognition algorithms against large volumes of network traffic. An initial configuration may work well, but networks are not static. Network traffic grows with business activity. Cybercriminals constantly develop new attacks and adapt to security countermeasures. IPSs are only as good as the policies they enforce, and so keeping policies up to date with the changing environment is crucial. Unless an IPS can help IT administrators understand their networks and endpoints and make recommendations about policies, the security device will either quickly degrade in its ability to meet core objectives or security specialists will be spending an inordinate amounts of time revising policies.

Adverse Impacts on Business Operations

The technical limitations of early IPS implementations led to several adverse business conditions. The most obvious was that senior management did not realize the value they thought they had invested in. The cost of evaluations, installations, configurations, tuning, and ongoing maintenance left them with information overload from excessive logging, network managers chasing false alarms, and concerns about protecting growing networks.

There was also an opportunity cost. The investment in early IPS technology still left network managers lacking valuable information about what was on their networks, such as vulnerabilities in devices hosted on the network. This reality leaves organizations to a range of threats from downtime to data loss—especially to the growing problem of targeted attacks.

Another unwelcome consequence is that IT staff time is spent on problems introduced by the IPS but that do not provide substantial value. For example, IPS administrators may spend time assessing irrelevant alerts and tuning rules to try to reduce false alarms and excessive logging. If nothing else, we have learned from these experiences that past evaluation criteria for IPSs have been insufficient.

More Effective IPS Evaluations

Rather than focus on limited evaluation criteria, such as the size of a rules database, a more pragmatic and ultimately effective approach is to evaluate an IPS in the broader terms of managing threats throughout the network—not just preventing isolated intrusions.

Targeted attacks on businesses are a growing problem. Cybercriminals can be more successful, from their perspective, by investing time in targeting and attacking a single business instead of unleashing generic malware into the wild and hoping it will eventually pay off. To combat this type of threat, network administrators should understand, in detail, the devices and services on their networks.

Data collected from IPS and network management systems can be leveraged to better understand threats to a network. What versions of operating systems (OSs) are running on servers? How many workstations are not patched to current levels? Are any devices out of compliance? These are the types of details that allow for preventive rather than reactive measures to ensure adequate security. Of course, evaluations are only the first step in realizing the benefits of intrusion prevention technology. Once IT knows what is typical on their network, only then can the IPS help them detect atypical behavior and thus help prevent threats.

More Effective Intrusion Prevention Deployments

Another set of factors to consider when choosing an IPS is the operational maintenance. Turnkey products, for example, reduce startup time. There is less configuration and less-demanding learning curves to deployment. Once an IPS is installed and operational, there will be maintenance issues as well.

A common maintenance task is optimizing rules for your particular environment. When comparing products, consider how an IPS supports rule tuning and optimization. This is an area that can potentially save significant staff time; especially important when one realizes that more senior professionals, not the junior staff, will be tuning IPS rules for the entire network. Again, IT can benefit from threats that have been detected. Look for detailed forensic tools that allow IT to drill down and truly understand the root cause of an attack. Then leverage the IPS to help recommend a change in policy that will prevent this attack from happening again.

Similarly, rules should be optimized to reduce false alarms and aid network administrators and security professional in focusing on top priorities. These include critical applications, enterprise databases, and the availability of network services.

Summary

Traditional IPS evaluations have not met expectations. Promises were made but not kept. Executives did not realize the return on investment they expected. IT professionals took on more work with overloaded logs and false alarms without the additional security they expected. This does not have to continue. Evaluation methods focused on broad security and business initiatives coupled can identify appropriate IPS solutions that are readily deployed and maintained.