Robert Drum
38 Schooner Ct
Richmond, CA 94804
510-237-1048
bobdrum@earthlink.net

1528 Words

IDS AND IPS PLACEMENT FOR NETWORK PROTECTION

by

Robert Drum, CISSP

26 March 2006

Introduction

This paper discusses the factors affecting proper placement of Intrusion Detection and Prevention System (IDS/IPS) sensors in computer networks. Differences between IDS and IPS capabilities and limitations of existing systems are explored. Given this background, appropriate deployment scenarios for IDS/IPS technologies are presented as well as some consequences of improper placement. Finally, security implications for network design and possible future enhancements to existing IDS/IPS systems are discussed.

IDS and IPS Technologies

Intrusion Detection and Intrusion Prevention Systems, IDS and IPS respectively, are mature network level defenses deployed in thousands of

computer networks worldwide. The basic difference between the two technologies lies in how they provide protection for network environments.

Intrusion Detection Systems, IDS, analyze network traffic and generate alerts when malicious activity is discovered. They are generally able to reset TCP connections by issuing specially crafted packets after an attack begins and some are even able to interface with firewall systems to re-write firewall rulesets on-the-fly. The limitation of Intrusion Detection Systems is that they cannot preempt network attacks because IDS sensors are based on packet sniffing technologies that only watch network traffic as it passes by.

Intrusion Prevention Systems, IPS, perform the same analysis as Intrusion Detection Systems but, because they are inserted in-line, between other network components, they can preempt malicious activity. In contrast to IDS sensors, network traffic flows through an IPS sensor not past it so the IPS sensor can pull or drop traffic from the wire.

This is the critical difference between IDS and IPS and it has implications for how both can be used. Because IPS sensors require traffic to flow through them, they can only be deployed at network choke points while IDS sensors can provide much broader network coverage.

Preliminary Information

Before discussing sensor placement, the target network should be analyzed and choke points identified. A choke point would be any point in a network where traffic is limited to a small number of connections. An example is usually a company's Internet boundary, where traffic crosses only a router and a firewall. The links between the router and firewall are perfect choke points and good places to consider placing IPS sensors.

Another consideration is high-value network assets. Business critical systems and infrastructure, such as server farms or databases, may warrant additional protection in the form of dedicated IPS or IDS sensors. Of course some of these assets can be protected by host-based IDS or IPS software agents in addition to, or instead of, targeted network level sensors.

Intrusion Prevention Sensor Placement

IPS sensors require network choke points; they are meant to be deployed between other network infrastructure components. An IPS sensor can only provide protection if traffic flows through it.

As we've seen, an Internet boundary is usually a good choke point, but there is another consideration in this case: do we position a sensor inside or outside the firewall? If we go outside, one sensor will protect the internal network and any DMZ networks behind the firewall. The downside is that the

sensor requires much more tuning to lower the noise level. Being outside the firewall means the sensor sees everything, even traffic the firewall would block. In this case, the IPS administrator needs to adjust the IPS policy or rule set so traffic that the firewall will block either doesn't get inspected by the IPS or the IPS doesn't generate alerts based on it.

This assumes that the administrator doesn't want to know about every inbound attack. In most corporate environments, this is true, but there are a few environments where it isn't, the individual administrator and their superiors must decide.

The flip side to this scenario is to place an IPS sensor inside, or behind, the firewall. Here, the firewall blocks traffic and therefore limits what the IPS needs to inspect, improving efficiency. The trade off is the number of sensors needed to provide the same level of protection as an externally placed sensor. Most commercially available sensors offer coverage for several physical network links in a single chassis or other hardware platform. Generally, the higher the number of links, the higher the cost.

Highly available networks add cost and complexity to both scenarios by increasing the number of physical links being protected. The decision of providing protection for the passive or fail-over side of a high availability lies with the system administrator and their superiors.

This discussion was specific to an Internet boundary but other likely choke points may exist. Many organizations maintain extranet connections to business partners that are consolidated on firewall or VPN protected networks. Placing an IPS sensor behind such a firewall or VPN concentrator protects one network from the other. In the case of VPN networks, care must be taken to inspect the un-encrypted side of the VPN tunnel.

There may even be choke points and boundaries within a network where IPS sensors can be deployed. Between departments or business units, or between users and critical systems like databases.

But what if a given network has no choke points? What this means is that flat networks are trouble for IPS sensors. But, in some cases, choke points can be created. Consider a switched network using one or no VLANs. On a single switch different ports can be assigned to different VLANs. Creating two VLANs and bridging them with an IPS sensor, creates a protected choke point . Network engineers will see this as an oddity and they are right but in a pinch, it works and allows different portions of the network to be protected from each other.

Another problem for IPS deployments is the wide-area network or WAN. IPS sensors can be used in wide-area networks but require positioning between distributed local area networks and the WAN cloud. This most likely translates to one IPS sensor at each remote location and one or more sensors at any central or

large sites. Obviously then IPS deployments in WAN environments can be expensive. I will leave one possibility up to the network engineers: in a hub-and-spoke WAN, it might be possible to leverage VLANS as discussed previously to get all traffic inspected by a single, centralized IPS sensor. This option is highly dependent on the given network infrastructure and also depends on all WAN traffic traversing the network through a single site.

Intrusion Detection Sensor Placement

As previously mentioned, Intrusion Detection System (IDS) sensors are more flexible and less capable than IPS sensors. Nonetheless, IDS sensors can be substituted for IPS sensors in all of the examples previously given and some of the same caveats apply, particularly when considering placement around firewalls. Importantly, though, IDS sensors forgo the need for in-line placement common to IPS sensors. IDS sensors can be connected to network taps or switch analysis ports, commonly known as SPAN ports. Both types of connections simply copy network traffic for presentation to and analysis by the IDS sensor. This means that IDS can provide security event detection with fewer sensors than IPS can, although the level of protection is far less.

For example, switched network backbones are ideal for IDS sensor deployment. Dependent on the amount of traffic being inspected, a few or perhaps even one IDS sensor can provide coverage for an entire network.

Actually, any switch that can enable an analysis port is a possible deployment site for an IDS sensor.

Implications for future IDS/IPS technologies

Clearly, existing IDS and IPS technologies have some limits, the need to protect at choke points only being chief among them. Aside from increases in processing speed, yielding the ability to inspect and protect more data per second, it seems that incorporating IDS and IPS technology into the network infrastructure is a logical next step. Some vendors are already providing something like this in the way of add-on modules or blades for existing switches.

But I think we will begin to see a hybridization of switch and security technologies in the next few years. A single device that appears to be a switch but has enough intelligence to perform a security analysis of not just every packet crossing the backplane but keep state on and watch every conversation, a session in network parlance. Such a device eliminates the need for separate IDS or IPS sensors sitting in the network and can conceivably protect system on adjoining ports from each other which is possible but cost prohibitive using today's technology.

These hybrid devices will be much more than just a switch with IPS. They will both require new technologies within the switch chassis and enable new

network architectures without. Whenever these devices arrive however, the need for them exists today.

Do note however, that the foregoing discussion does not mention firewalls. The merger of firewalls and IPS/IDS technologies isn't necessarily logical. Firewalls are designed for very rapid inspection of packet headers so they can make very rapid decisions about passing traffic. Intrusion Detection and Prevention Systems are designed to delve far deeper into packets and entire network sessions. I think it will be many years before we see network devices that can effectively deal with both of these jobs.