Intrusion Detection and Prevention Systems Simplified

Arthur J. Wyatt

East Carolina University

WWW.INFOSECWRITERS.COM

Abstract

This paper attempts to cover and discuss several aspects of Intrusion Prevention Systems and Intrusion Detection Systems. This paper attempts to do this is simple and basic language. Both systems are explained and defined according to the National Institute of Standards and Technology. Following that several techniques that can be used to install or implement them are described. The techniques described are hubs, port mirroring, test access points, and inline. During the discussion of each there are figures to depict and assisting in conveying how each implementation works. How each works along with the security or performance issues is discussed. The last thing talked about in this paper network segmentation and how Intrusion Prevention Systems and Intrusion Detection Systems could be used in conjunction to layer security and enforce network use and security policies.

WWW.INFOSECWRITERS.COM

Intruduction

One of the biggest challenges Network Administrators face today is being able to see what is going on in their own network. It is no longer enough to know that a service or machine is up and running. On a single workstation it may be easy enough for one person to have a good idea of what is happening, but when that changes to several machines, sometimes several thousand machines, it becomes increasingly difficult to have a clear picture of who is on what machine and what they are doing. This does not even address the issue of holding an end user accountable for misusing company resources or identifying a security threat that has breached the firewall. The solution to this is to use network monitoring tools to increase the visibility of the network. These tools can be used to monitor network performance or help secure it through Intrusion Detection Systems and Intrusion Prevention Systems. This paper will discuss some techniques on how and where to implement the network monitoring devices, and will conclude with suggestions and a conclusion.

Intrusion Detection Systems

First, it is important to understand what an Intrusion Detection System (IDS) and Intrusion Prevention System (IPS) are. The National Institute of Standards and Technology (NIST) has three definitions for IDS and one for IPS. According to NIST IDS is defined as follows:

> Intrusion Detection Systems (IDS) – Hardware or software product that gathers and analyzes information from various areas within a computer or a network to identify possible security breaches, which include both intrusions (attacks from outside the

WWW.INFOSECWRITERS.COM

organizations) and misuse (attacks from within the organizations.) SOURCE: CNSSI-4009

Intrusion Detection Systems (IDS) – (Host-Based) IDSs which operate on information collected from within an individual computer system. This vantage point allows host-based IDSs to determine exactly which processes and user accounts are involved in a particular attack on the Operating System. Furthermore, unlike network-based IDSs, host-based IDSs can more readily "see" the intended outcome of an attempted attack, because they can directly access and monitor the data files and system processes usually targeted by attacks. SOURCE: SP 800-36; CNSSI-4009 NIST IR 7298

Intrusion Detection Systems (IDS) – (Network-Based) IDSs which detect attacks by capturing and analyzing network packets. Listening on a network segment or switch, one network based IDS can monitor the network traffic affecting multiple hosts that are connected to the network segment. (Kissel, R.)

IDS, broadly, is any device or application that can monitor and analyze data looking for specific events or type of events. Then when one of the specific events occur the IDS can generate an alert and email the System, Network, or Security Administrator. NIST, however, also makes the important distinction that there are primarily two kinds of IDSs: host or network based.

Host-based IDS function either as a standalone on a single machine or in a server-agent compacity. Both operate in the same basic way. They monitor for changes and actions on a host looking for events to trigger an alert. Some of the characteristics that the IDS could monitor include "wired and wireless network traffic, system logs, running processes, file access and

WWW.INFOSECWRITERS.COM

modification, and system and application configuration changes (Kent, K.)." On a standalone system the IDS would monitor the events and log them on the same machine. This works great for small networks or even a home user; but if an attacker gained access to the machine it would be theoretically possible for the malicious user to delete the logs and erase all trace of the attack before the system administrator has a chance to view or check the logs. There are a few ways to combat this. The first is to setup any form of remote logging. The most basic of which would be to have the alerts emailed to the administrator as they occur. A better way might be to set up a server-agent IDS system. In a server-agent IDS the agents are all hosts that need to be monitored. The agents will need to have the necessary hard or software installed and configured. Part of this setup would include designating a 'server'. The server could be any of the hosts or another machine entirely, which is suggested. Information would be sent to the server machine to be processed. This would free up CPU, MEM, and storage resources for the agents. Once the server receives the information from the agents it would then use its own resources to look for events and generate alerts logging them on the server. This would create remote logging and give the system administrator a central place to view and monitor all of the configured agents.

Unlike Host-based IDSs, which monitors system process and resources, IDS Network-based IDS monitor exclusively network traffic. It does this by grabbing the network packets while they are in transit. In order for the IDS to be able to sniff or grab every packet, even those which the IDS machine is not the destination, the network interface card or NIC would need to put in to promiscuous mode. While in this mode the IDS would be able to grab all network packets indiscriminately. Network-based IDSs perform most of their analysis based on the application layer [e.g. File transfer protocol, FTP), transport layer [e.g. TCP, UDP, ICMP], and network layer [e.g. ipaddress]; but sometimes it can also perform limited analysis based on the

WWW.INFOSECWRITERS.COM

hardware layer(Kent, K.). There are two different types of Network-based IDSs installations, inline or passive. In an inline installation the Network-based IDS is directly in the flow of the network traffic. This forces all traffic to go through the IDS. An example of this is shown in Figure 1.

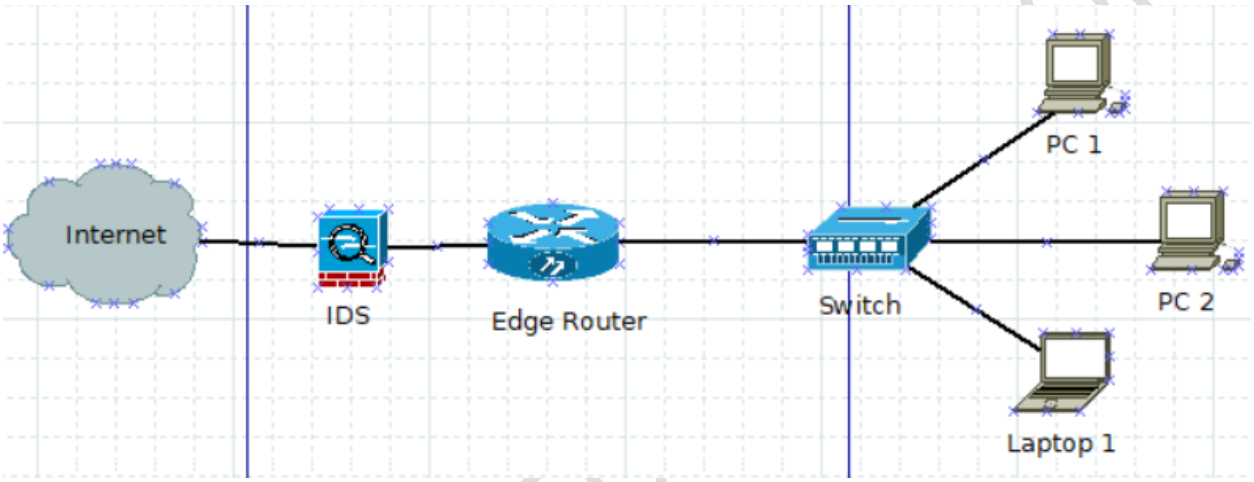

Figure 1.

In Figure 1 the IDS is placed inline between the edge router and the internet. Because of this, all the traffic from the internet to the inner network for PC 1, PC2, and Laptop 1 and vice versa will have to go through the IDS. This can cause a bottleneck to form and if the IDS for any reason goes down or malfunctions then all network traffic destined to go to the internet or from the internet to the inner network will fail, effectively shutting down the network. Because of this, it is typically a better idea to place the IDS in passively.

There are a number of ways that a IDS can be placed in a network to passively collect data, but this paper will only cover TAPs, port mirroring, and hubs. Hubs are devices like

switches, but instead of selectively sending packets to only the intended recipient a hub will broadcast all packets out of all connected ports regardless of which host is connected to them. Because of this, an IDS connected to a hub will receive a copy every packet that passes through the hub allowing the IDS to monitor the traffic, analyze it, and generate alerts if necessary. Figure 2 shows an example of this type of setup.

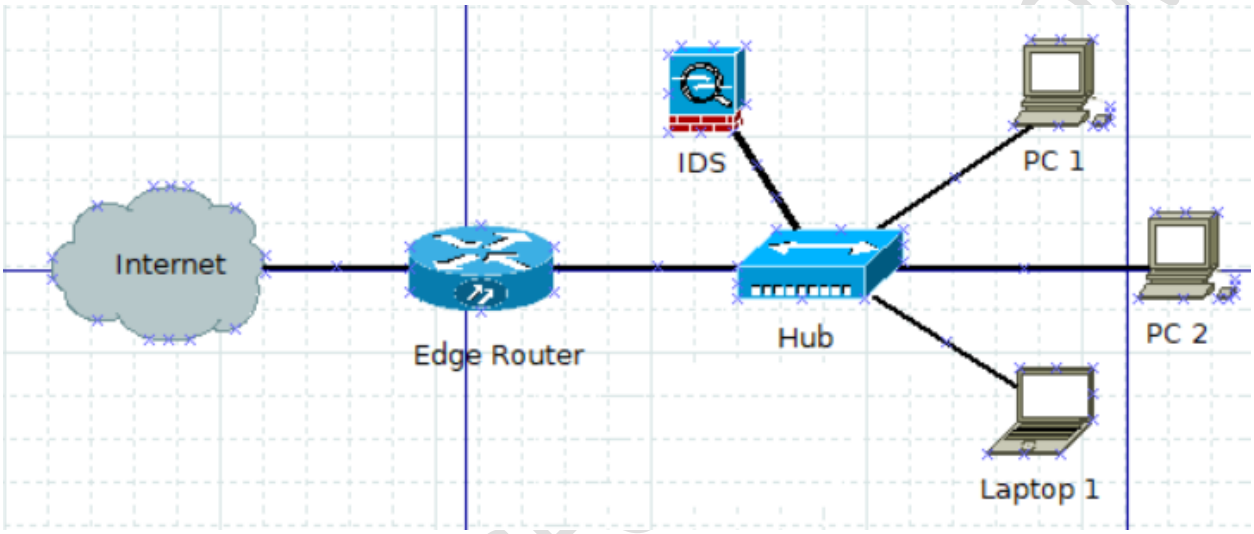

Figure 2.

Figure 2 shows an IDS that is connected to a hub on the inner network. As mentioned earlier, This allows the IDS to receive a copy of every packet that passes through the hub. However, this also allow any other device connected to the hub to do the same. If a malicious user was aware of the hubs existence the user could connect another device to the hub that acts similarly to the IDS and passively collects data and analyses it. The goal of this malicious user would then be able to use any information learned or directly stolen to perform other attacks on the network and because the devices would also be in passive mode and not actively sending traffic through the hub there would be limited ways, if any, for the IDS to detect the additional device and alert the system administrators. For this reason, hubs are generally considered to be

insecure and only used when absolutely required. The other ways mentioned in this paper; port mirroring, and TAPs, are suggested instead because of their increased security.

Port mirroring is the generic term used to describe the protocols used and process of sending all the data that is destined for one port to another port as well. This is also commonly known by what Cisco's term Switched Port Analyzer or SPAN (Rouse, M.). Using port mirroring or a SPAN port a system administrator can achieve the same result as if a hub were used. An example of this setup is shown in figure 3.

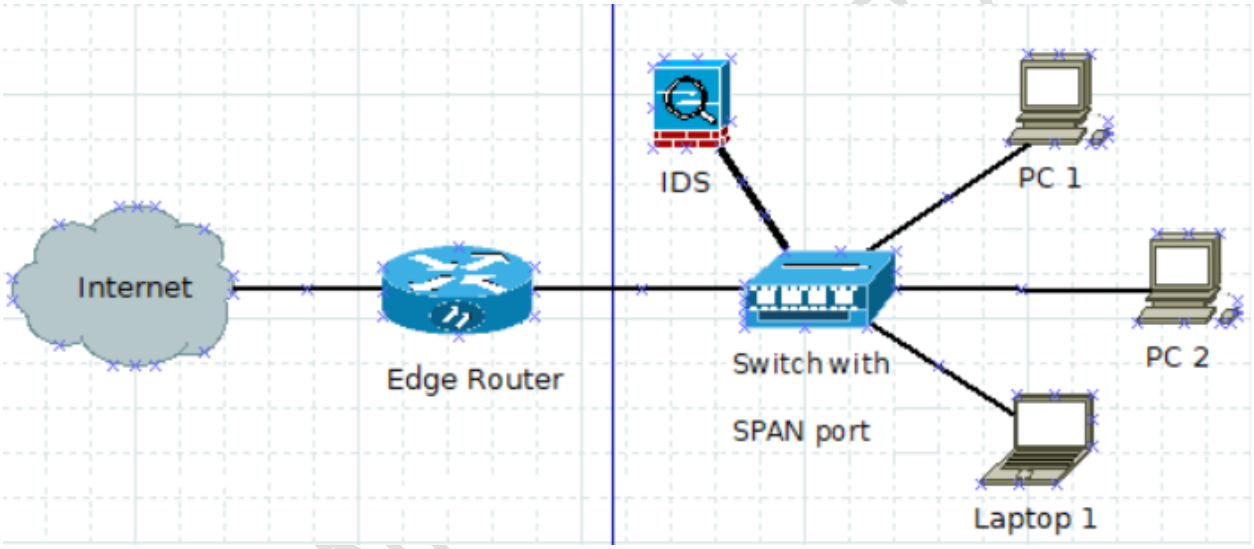

Figure 3.

The network topology for figures 2 and 3 are similar and achieve a similar result. Both the hub and port mirroring configurations will allow the IDS to monitor traffic and generate alerts as needed. However, there is one significant difference. It was mentioned earlier that the use of a hub on the network could potentially be a security risk and compromise the network because the hub would indiscriminately send all data to all ports. This potentially would allow a hacker to gain access to data they would not otherwise have access too. Using port mirroring or

SPAN could be a solution to this issue. Using a switch instead of a hub would prevent another user from connecting directly to the device and passively gathering network traffic. Using a switch could also pose problems as well. Because not all switches were created equal the switch would first need to be capable of using port mirroring. Then it would need to be configured to do so. Once configuring the switch would then send a copy of packets to the SPAN port connected to the IDS. This additional work load would require more resources from the switch and if it is over-worked could cause other issues.

The last implementation of passive monitoring covered in this paper is TAPs. Test Access points or TAPs are devices that can be placed in between two nodes (Rouse, M., & McGilicuddy, S.). Quite literally a TAP can be placed in the middle of a network wire or cable connecting any two devices (e.g. a router and switch). A TAP serves the same purpose as port mirroring in that it will mirror all traffic and relay it along another cable to a monitoring device. This requires additional hardware over using a SPAN capable switch, but requires fewer resources because the TAP device is independent. An example of such a configuration is shown in figure 4.

INFOSECWRITERS.COM

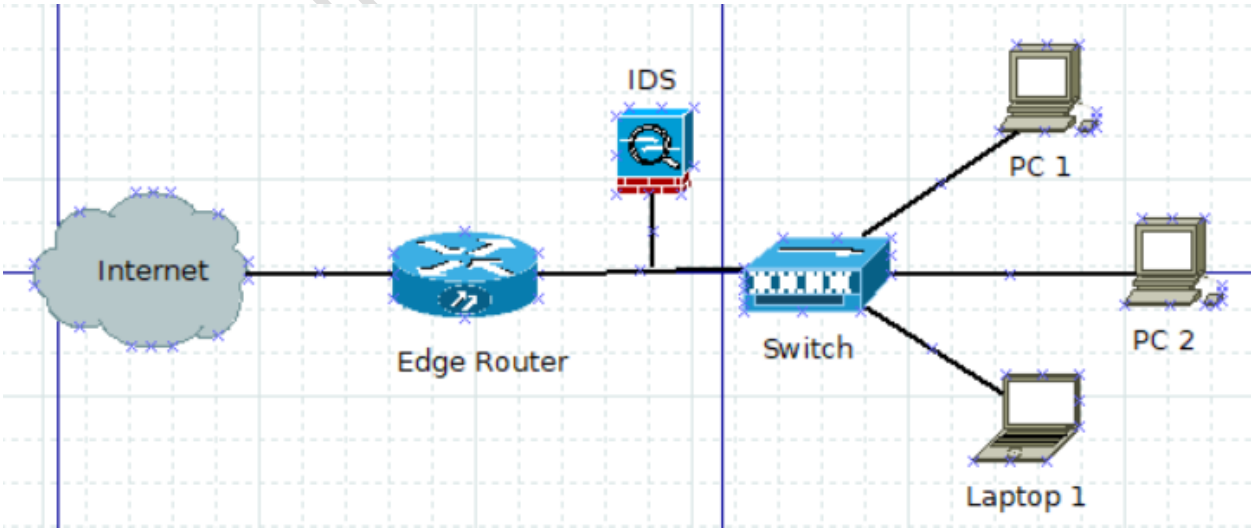

Figure 4.

In figure 4 the IDS has been installed via a TAP between the edge router and switch. This type of configuration works great if an inline installation is not wanted and/or a SPAN capable switch is not available. Even if there is a SPAN capable switch around it may be preferred to use a TAP because of the lessened strained on the switch's resources.

Intrusion Prevention Systems

IPSs are fundamentally different than IDSs. NIST defines IPSs as "System(s) which can detect an intrusive activity and can also attempt to stop the activity, ideally before it reaches its targets" (Kissel, R.). The key difference between IDSs and IPSs is that IDSs passively monitor traffic, but take no preventative action to stop an ongoing attack, While, IPSs will passively monitor traffic, but in addition will take action and actively try to prevent any traffic it determines to be malicious. An example of using an IPS in a network is shown below in figure 5.

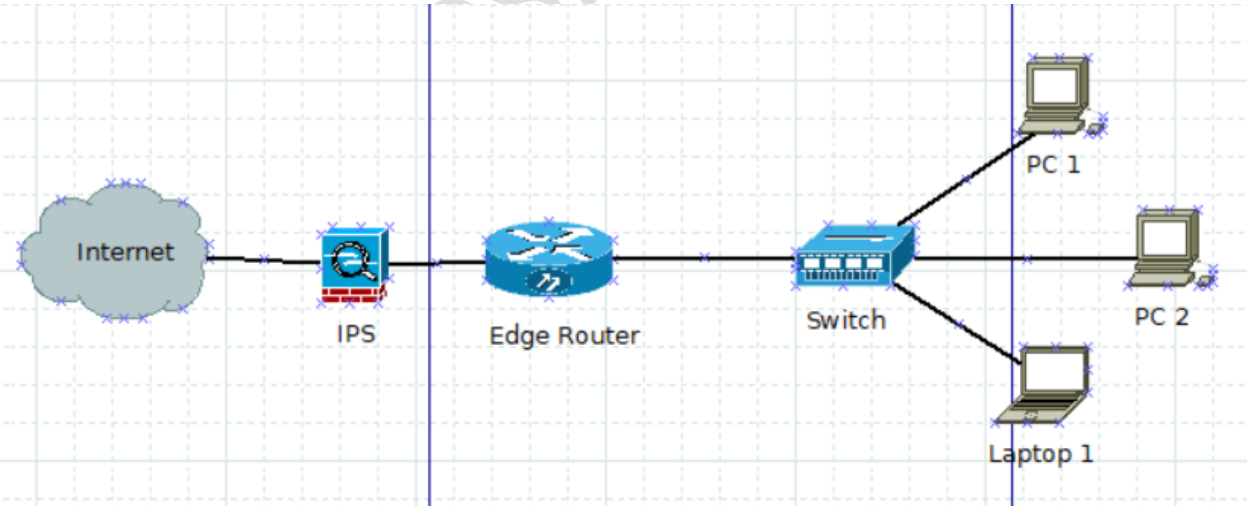

Figure 5.

IPSs use many of the same techniques to recognize harmful activity on the network as IDSs do, but because an IPS' purpose is to prevent the attack they should be installed inline as figure 5 shows. Being installed inline the IPS can drop or reroute packets as they pass through the system. If the IPS was setup in one of the other passive ways, such as a TAP or SPAN port, then any manipulation of the packets by the IPS would not have any actual effect. For example, If the IPS was configured on a switch's SPAN port to receive mirrors of all packets that pass through it, then once dangerous activity is matched it drops the packets. This would only have an effect on the mirrored packets and would have no effect on the original traffic sent to the victim machine. Thusly, Installing an IPS in a passive way would not prevent the packets from reaching the victim and would effectively have the IPS function in the same capacity as an IDS.

IPSs can also be used to reduce the strain on the entire network. Because IPSs can be configured to simply drop packets an IPS can stop network devices, like routers and switches, from ever having to process them. A basic implementation of this strategy is shown in figure 5. In figure 5 the IPS is positioned between the internet and the edge router. If an attacker tried to probe or attack the network the IPS could identify the intrusive behavior and drop the packets. This both defends against the attack and prevents the router from have to process what to do with it. This would free up resources and benefit a healthier network.

WWW.INFOSECWRITERS.COM

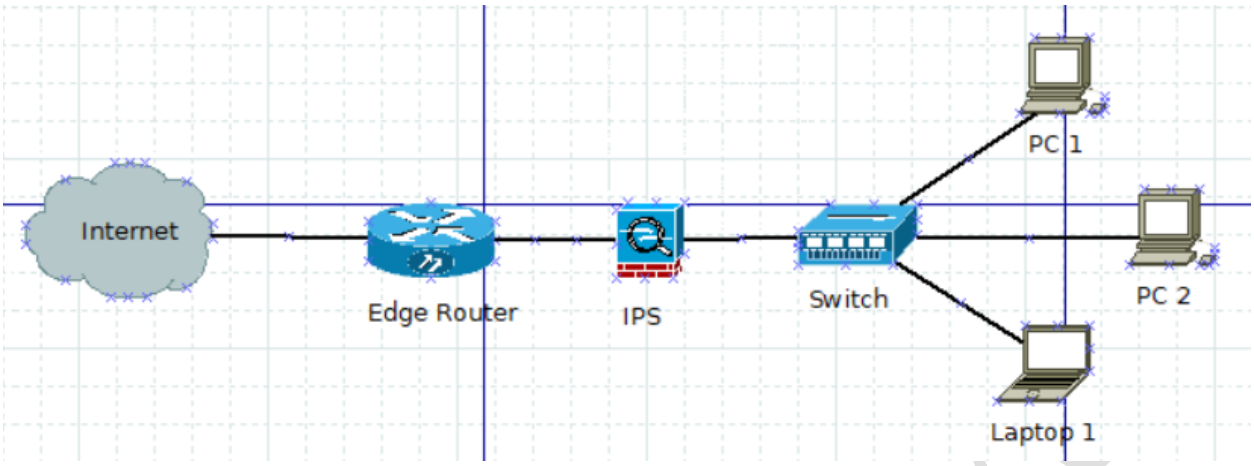

Figure 6.

Conversely if the IPS was placed on the other side of the router as shown in figure 6 then all the traffic would still have to be processed by the edge router before the IPS could filter it. In this example attacks destined for PC 1, PC 2, or Laptop 1 could still be prevented as long as the attacks were recognized by the IPS device. However, in this second example the edge router would be under increased strain due to the fact that all the traffic would have to processed by the router before the IPS would be given the opportunity to act upon it.

Detection Methodologies

IPSs and IDSs use the many of same technologies to detect known or potential threats. There are 3 primary ones. They are signature-based, anomaly-based, and stateful protocol analysis (Kent, K.). Each of these function differently and are used for different reasons. Some IPSs or IDSs may only use 1 of the 3 while others will use any combination of them. This next section of the paper will discuss at a basic level how they work and give a few examples.

Signature-Based Detection

WWW.INFOSECWRITERS.COM

Signature-based detection works by looking for defined patterns called signatures. These signatures can be predefined by default and additional ones can be defined by the end user, security, or system administrator. This detection method is very literal and will only catch events that match exactly. Even something as simple as changing the name of a file can allow something to slip through signature-based detection. Because of this, though it is very good at detecting known threats. If there are premade threats or mass rolled out attacks they can be defined in the IPS or IDS and they will catch it. This type of detection is also efficient at detecting company security, or other, policies. A few examples of signatures are below:

- A telnet attempt with a username of "root", which is a violation of an organization's security policy
- An e-mail with a subject of "Free pictures!" and an attachment filename of "freepics.exe", which are characteristics of a known form of malware
- An operating system log entry with a status code value of 645, which indicates that the host's auditing has been disabled (Kent, K.)

Signature detection is pretty basic. All it does is compare something (e.g. packets or logs) with a comparison string. If there is a match then X action is taken. Also, Signatures do not look at the overall picture. Once a comparison has been made it then immediately forgets about it and moves on the next. Because of this, signature detections do not notice correlations over time.

Anomaly-Based Detection

The second primary type of detection is anomaly-Based. Unlike signature-based which uses comparative strings, anomaly-based detection compares behavior. Initially and IPS or IDS

WWW.INFOSECWRITERS.COM

would passive monitor all traffic without generating alerts to get a 'feel' of what is normal network behavior. This is often called the training period. An example would be if the device monitored for a week and found that on the average work day around 20 emails were sent to and from the network per half hour. After that the IPS or IDS would generate an alert anytime the number of emails per half hour was abnormally above 20. If, for a period of time, the number of emails per half hour was 100 or even 1000 that could be the symptom of a compromised machine. For this reason, the anomaly detection method is good at catching previously unknown or zero-day attacks.

Anomaly-based detection is not perfect however. If the baseline or average is static it would need to be updated from time to time because of the changing needs of the network. What is average for one week may not be the same in a few months or years. If the time between recreating baselines is too long the risk of false positives increases. If the profile is dynamic it would be constantly updated and changes in the network's needs would automatically be incorporated. This would make maintaining the IPS or IDS easier and more efficient, but dynamic profiles are susceptible to evasion attempts. Because dynamic is constantly updated a malicious user or attacker could perform attacks over time and in small increments. By increasing the scale and frequency over time the dynamic learning would learn that these attacks were normal and incorporate them in to the profile.

Another issue is that often times there are activities such as maintenance or backups that only have a few times a year. With either static or dynamin is it difficult to incorporate these events in to the IPS or IDS without making the profile inaccurate for rest of the time. Often times these events will trigger a false positive alert and it is up to the person responsible for monitoring the alerts to recognize this.

WWW.INFOSECWRITERS.COM

Stateful Protocol Analysis

The last of the primary detection methods is stateful protocol analysis. Stateful protocol analysis works by comparing the protocol state with a predefined profile that determines if activity is suspicious, malicious, or benign. A state is the current condition that a protocol is in. for example, when a user accesses a File Transfer Protocol (FTP) server but does not provide a valid user account and password combination then the FTP is in an unauthenticated state. In this state typically, a user would have access to a small set of commands such as the 'help' command (Kent, K.). If the user attempted to use any other command then this could be caught by the IPS or IDS and generate an action or alert. Conversely, if the user did provide a valid account name and password combination then they FTP would be in an authenticated state. This change in state would cause many things that would be suspicious in an unauthenticated state to now be considered begin.

Stateful protocol analysis can also monitor and analyze the use of commands and the sequence of commands used. Some commands usually only used before or after another related command. If it is attempted to execute these commands out of order then that could be a sign of misuse or intrusion. Also, if a command typically accepts arguments of a particular type and length then something other could be a bad sign. For example:

> If a command typically has a username argument, and usernames have a maximum length of 20 characters, then an argument with a length of 1000 characters is suspicious. If the large argument contains binary data, then it is even more suspicious(Kent, K.).

Vendors define the characteristics of their protocols and the corresponding profiles. Because of this sometimes a profile is not comprehensive or may have never been created, especially for

WWW.INFOSECWRITERS.COM

proprietary protocols. Another primary drawback is that there is substantial overhead and can be resource intensive for a IPS or IDS to track multiple sessions. If there are too many it may become impossible for the monitoring devices to accurately and effectively monitor all instances. Also, if the protocol is used in a way that is not defined by the vendor there could be conflicts.

Network Design

This section will briefly discuss using IPS and IDS together along with network segmentation to achieve a more secure network. So far, this paper has covered some basic examples of where and how they can be implemented, but perhaps a better idea would be to use them in conjunction with each other. An IPS configured to be very strict could potentially do more harm than good. Because an IPS is designed to take action when a match to a signature rule or a deviation from normal activity in behavior based is found it should try to take measures to stop the 'attack'. However, if this traffic is a false positive and in fact legitimate traffic this would negatively affect the end user experience and could go so far as to make the network unusable. Conversely, even if an IDS is configured to be strict there should be no negative impact on the network because no action or preventative measure is taken. Therefore, it could be a more intelligent to use a conservatively configured IPS with a strictly configured IDS. This would allow a System Administer to monitor the alerts generated by the strict IDS and if there are less than the acceptable number of false positives then that rule can be transferred to the IPS. This way experimental rules are tested for functionality before being used in the IPS and potentially harming the network experience. A simple network design using both an IPS or IDS could look like the one depicted in figure 7.

WWW.INFOSECWRITERS.COM

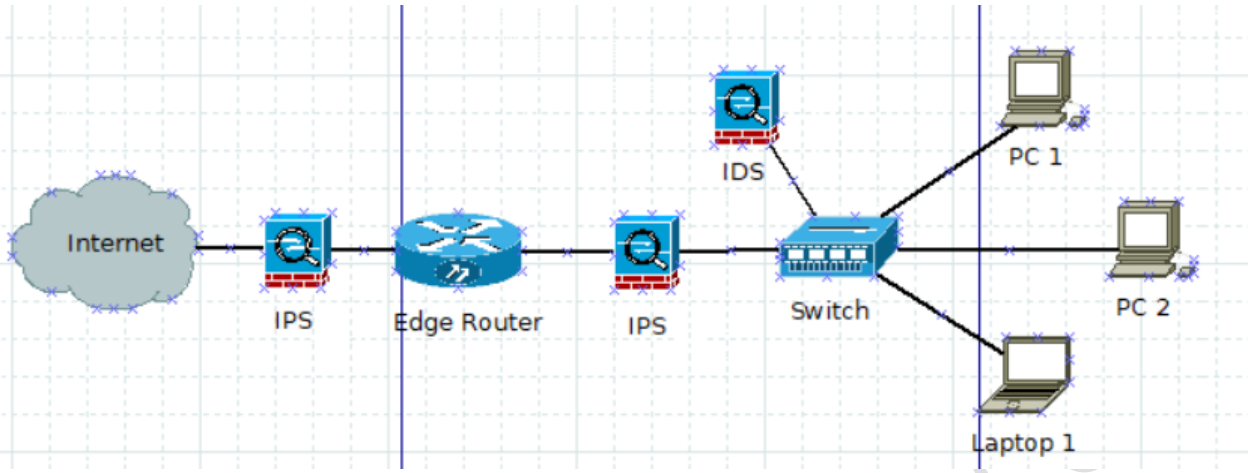

Figure 7.

In figure 7 there are two separate IPSs is installed inline. The first is between the edge router and the internet. This way any traffic that is matched can be dropped before it reaches the router and will save some resources. However, because of Network Address Translation it may be difficult to use this IPS to match traffic to rules based on the inner network topology. To solve this the second is installed between the switch and the edge router. From here the IPS can match more granular rules that is matched against the end devices and because of its location it can drop traffic before the hosts have to waste resources on processing it. The IDS is installed on a SPAN port of the switch on the inner network. From the SPAN port the IDS can get a copy of all the traffic sent through that switch and test it verse its own rules to generate alerts if needed. Experimental rules can be set in this IDS and tested to see if they work as intended. If they do not the rules could be modified until that work as expected or deleted if deemed not necessary. More importantly, if the rule works well it can be moved over to the inner IPS, the one between the edge router and switch.

Bigger networks will typically require more than a few end devices and workstations and it does not always make sense to keep all of these on the same inner network. A strategy that is

WWW.INFOSECWRITERS.COM

often used to help isolate and separate them is to use network segmentation. Network segmentation can be loosely defined as isolating access to end devices based on logical groupings, minimal access needed, and trust levels. An example of using network segmentation is shown in figure 8.

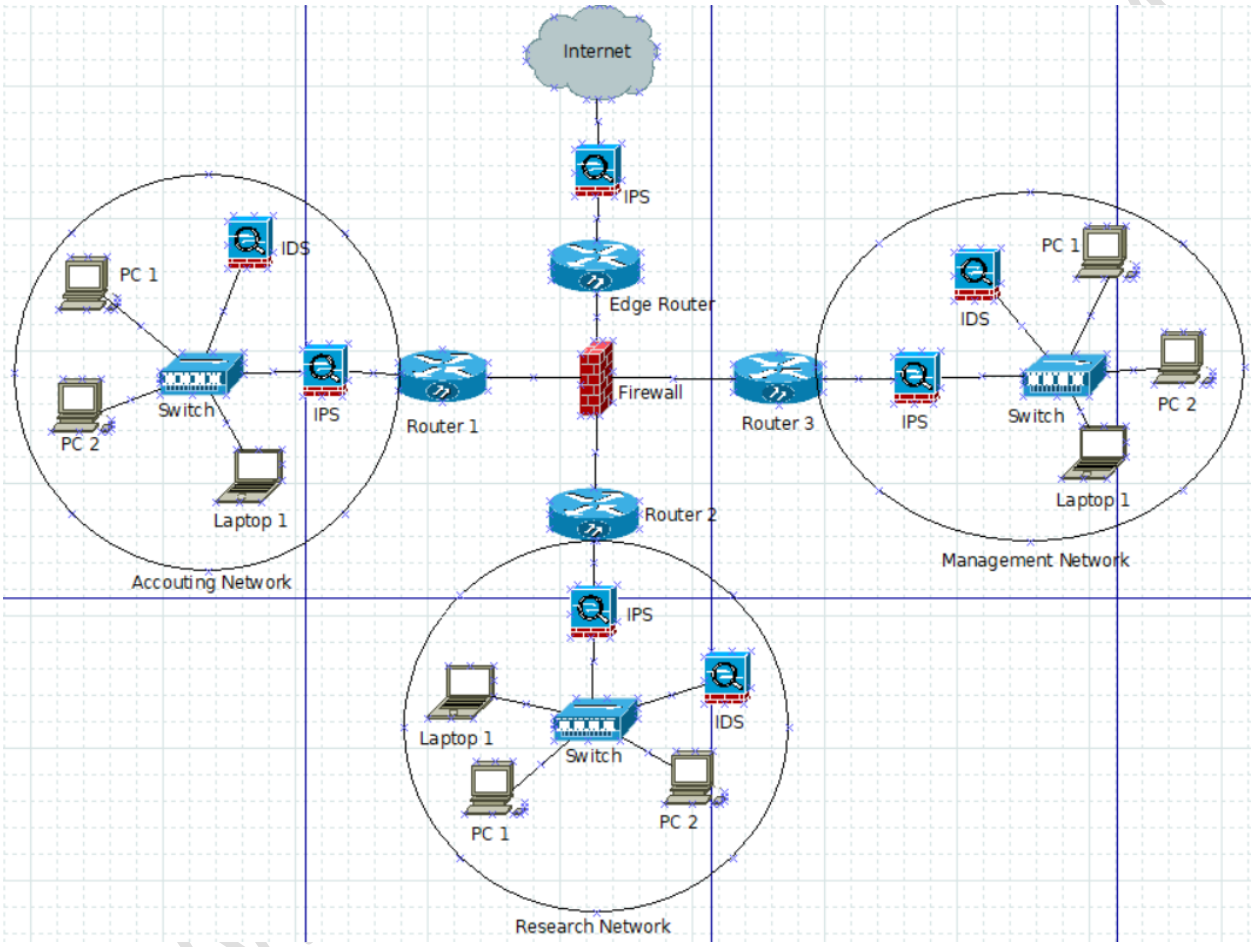

Figure 8.

In figure 8 there are 3 different networks: accounting, Research, and management. Each of these is logically grouped where every device or employee would be placed in the respective network. Individuals and devices in accounting would be placed in the accounting network, research in research and management in management. Next policies would need to be created to

define what is and is not allowed. These policies should try to take in to account what each needs access too and only grant that while simultaneously denying access to everything thing else. Questions like "should accounting have access to the management network?" should asked. In most cases the three networks would not need any access to the other 2, except for management which may need to be granted access to research, but research would probability not be granted access to the management network.

In the topology shown in figure 8 each network has an IPS in between the switch and router pointing to the firewall. This IPS would be used to enforce the network and security polices created. For example, if someone in the accounting network tried to access the management network the IPS should drop the traffic and generate an alert so that, if necessary, whatever or whoever is reasonable can be investigated for accidental mishap or malicious intent. Each Network also is configured with an IDS on the on the SPAN port of the switch. This would be used to reinforce the policy on a stricter scale than the IPS would be able to without adversely affecting the network user experience.

Conclusion

In conclusion, IPS and IDS are a fundamental part of any network security setup. IPSs can be used to prevent an attack or misuse by dropping or rerouting the traffic. However, to prevent the IPS from affecting legitimate it should be configured conservatively with well tested rules. An IDS can be configured to be stricter because it does not actively affect traffic and the administrator would be responsible for looking over the alerts generated to determine if they are true or false positives. If the rules set in the IDS are thoroughly tested than can then be moved to the IPS. Using IPSs and IDSs in conjunction with each other would greatly increase the security

WWW.INFOSECWRITERS.COM

of a network when properly configured and used. Lastly, IPSs and IDSs can be used to add an additional layer of security to other strategies such as segmentation.

WWW.INFOSECWRITERS.COM

REFERENCES

Asiwe, V., & Dowland, P. (n.d.). Implementing Network Monitoring Tools. Retrieved March 25, 2018, from https://www.cscan.org/download/?id=383

Carlo, C. D. (2003, September 25). Intrusion detection evasion: How Attackers get past the burglar alarm. Retrieved March 25, 2018, from https://www.sans.org/reading-room/whitepapers/detection/intrusion-detection-evasion-attackers-burglar-alarm-1284

Kent, K. (. A., Mell, P., & National Institute of Standards and Technology (U.S.). (2007). Guide to intrusion detection and prevention systems (IDPS): Recommendations of the national institute of standards and technology. Gaithersburg, MD: U.S. Dept. of Commerce, Technology Administration, National Institute of Standards and Technology.

Kissel, R., & National Institute of Standards and Technology (U.S.). (2011). Glossary of key information security terms (Revision 1. ed.). Gaithersburg, Md.: U.S. Dept. of Commerce, National Institute of Standards and Technology.

Pappas, N. (2008, April 2). Network IDS & IPS Deployment Strategies. Retrieved March 25, 2018, from https://www.sans.org/reading-room/whitepapers/intrusion/network-ids-ips-deployment-strategies-2143

R. (2017, December 12). The Pros & Cons of Intrusion Detection Systems. Retrieved March 25, 2018, from https://blog.rapid7.com/2017/01/11/the-pros-cons-of-intrusion-detection-systems

Rødfoss, J. T. (2011, May 24). Retrieved March 25, 2018, from https://www.duo.uio.no/bitstream/handle/10852/8951/Rodfoss.pdf

WWW.INFOSECWRITERS.COM

Rouse, M. (2014, March). What is port mirroring (roving analysis port)? - Definition from WhatIs.com. Retrieved March 25, 2018, from https://searchnetworking.techtarget.com/definition/port-mirroring

Rouse, M., & McGilicuddy, S. (2013, May). What is network tap? - Definition from WhatIs.com. Retrieved March 25, 2018, from https://searchnetworking.techtarget.com/definition/Network-tap

Sy, B. K. (2009). Integrating intrusion alert information to aid forensic explanation: An analytical intrusion detection framework for distributive IDS. Information Fusion, 10(4), 325-341. doi:10.1016/j.inffus.2009.01.001

WWW.INFOSECWRITERS.COM