Network Security

Professor Gaurav S. Kasbekar
Department of Electrical Engineering
Indian Institute of Technology, Bombay

Week - 10

Lecture - 55

Firewalls and Intrusion Detection Systems: Part 5

Hello, recall that in the previous several lectures, we discussed firewalls, and we also discussed the different types of firewalls. We will now discuss intrusion detection systems in this lecture and the next few lectures. So, what is the need for intrusion detection systems? Recall that a packet filter, whether it is traditional or stateful, inspects different header fields, including IP, TCP, UDP, and ICMP header fields in the network layer header, transport layer header, and so on. So, by inspecting these fields, the packet filter decides whether to let a packet pass or to block it.

However, to detect several kinds of attacks, it is not sufficient to just inspect the header fields. We need to perform what is known as deep packet inspection. That is, we need to look at the application data that packets carry in addition to the header fields. So, it is not sufficient to just examine the header fields. We need to examine the header fields and also the application data that the packets carry.

Let's look at an example of such an attack for which we need to perform deep packet inspection to detect that attack. Consider packets carrying viruses or worms. So, this virus or worm data is in the application part of the packet, that is, in the payload of the packet. So, to detect that the packet is carrying viruses or worms, we need to look at the application data in the packet, and that will tell us that there is some malicious information in the packet. So, looking at the header fields will not tell us the packet is malicious.

Now recall that application gateways, which we discussed earlier, perform deep packet inspection. For example, we looked at an example of a Telnet gateway. It examines the application data in the packet and finds out which user is trying to Telnet. So, application gateways perform deep packet inspection. But they only do this for a single application, for example, Telnet, HTTP, FTP, and so on.

So, there is a separate application gateway for each application. Hence, there is a need for another kind of device that examines the headers of all the packets passing through it, just like a packet filter. But in addition, it also examines the application data contained in the packets. That is, it performs deep packet inspection. So, firewalls perform a useful role, but in addition to firewalls, we also need another kind of device that examines the headers of packets as well as examines the application data that is there in the packets. So, such a device is known as an intrusion detection system.

So, when an intrusion detection system observes a suspicious packet or a suspicious series of packets, the device may either prevent the packets from passing through or let them pass but send alerts. So, a device that generates alerts when it observes potentially malicious traffic is called an intrusion detection system. And a device that filters out suspicious traffic is called an intrusion prevention system. So, the difficult part is not whether to generate alerts or to filter out suspicious traffic. So, the challenging part is to detect malicious traffic.

We study IDS and IPS together because the challenging part is to detect suspicious traffic. Once suspicious traffic has been detected, either sending alerts or dropping packets is straightforward. Hence, we study IDS and IPS together. So, in both these kinds of devices, the challenging part is to detect malicious traffic. Once this detection has been done, if it is an IDS, then it will generate alerts, and if it is an IPS, then it will filter out suspicious traffic.

So, we collectively refer to IDS as well as IPS as IDS, intrusion detection systems. So, let's discuss some example attacks that can be detected using intrusion detection systems. One class of attacks that can be detected using IDS is network scanning. This includes discovery of hosts, services, and vulnerabilities on a computer network by sending probe packets into the network and analyzing the responses. So, an attacker wants to scan the network and find out which hosts are there in the network, which services the hosts are running and what are the vulnerabilities on the different hosts.

So, it sends a series of probe packets into the network, and then the hosts of the network respond to these probe packets, and by analyzing these responses, the attacker gets to know which are the host services and vulnerabilities on the network. So, network scanning can be performed using software such as "nmap". Here's an example of network scanning. One is host discovery, that is, identifying which hosts are there in the network. The host discovery can be done as follows.

The attacker can list the hosts that respond to TCP and/or ICMP requests. So, the attacker can send some TCP and/or ICMP requests into the network, and when it sends such requests to a particular IP address, if it gets a response, then it means that there's a host at that IP address. Then another example of network scanning is port scanning. The attacker lists the open ports on the target host. So, again, in port scanning, the attacker sends a series of packets destined to different ports, and if it gets a response, it means that there is some process at that particular port of the target host.

So, in this way, it can find out at which ports there are processes running. Then another example of network scanning is operating system and hardware detection. The attacker can determine the operating system and hardware that is deployed in host. So, again the attacker can send some probe request packets into the network, and by analyzing the responses, it can find out which operating system and hardware is deployed in the host. Then another example of network scanning is network vulnerability scanning.

So, an example is a tool known as SAINT, which stands for Security Administrator's Integrated Network Tool. It's a computer software that detects the TCP and UDP services running on every host of a network. So, for every service that it finds running, it sends a series of probe packets, which are designed to detect weaknesses that could allow an attacker to gain unauthorized access, launch a denial of service attack, or gain confidential information. In this way, an attacker can analyze the vulnerabilities in a network. So, later it can use the knowledge of these vulnerabilities to attack the network, possibly sending some well-designed packets to a host so that the vulnerability causes the host to crash.

So, that's one example where the detected vulnerability is exploited. Then, another example is OS vulnerability attacks. Then, application vulnerability attacks; these are other examples of attacks that can be detected using intrusion detection systems. Then, the injection of malware, for example, worms and viruses, into the hosts of the network. So, an IDS can analyze packets, perform deep packet inspection, and if some of the data in the packet is malicious, it may indicate that there is some worm or virus in the packet. So, the injection of malware into the hosts of the network can be detected using IDSs. Then, application-layer denial-of-service or distributed denial-of-service attacks. That is another example of an attack that can be detected using intrusion detection systems. An example is when an attacker sends a large number of requests to log into an online account, such as a Gmail account.

So, this is an example of an application-layer DoS or DDoS attack. So, the way this hampers the victim is that a lot of server resources are consumed in the process of loading the relevant user data from a database, checking login credentials, and sending a response containing the requested webpage. So, a lot of resources are consumed in these processes. So, that's how it adversely affects the victims. So, these are all examples of attacks that can be detected using IDSs.

So, we see that there is a wide range of possible attacks that can be detected using intrusion detection systems. Now, the question is: how does a network that contains intrusion detection systems look like? So, we'll consider an example network architecture that includes IDSs. An organization may deploy one or more IDS sensors in its organization's network. Typically there are multiple IDS sensors, and each IDS sensor analyzes a part of the traffic that flows through the network.

This figure shows an organization that has three IDS sensors. So, this symbol denotes an IDS sensor. So, this is the public internet, and this is the organization's network. The organization's network is divided into two parts. One is a low-security zone called the demilitarized zone, and the other is a high-security zone.

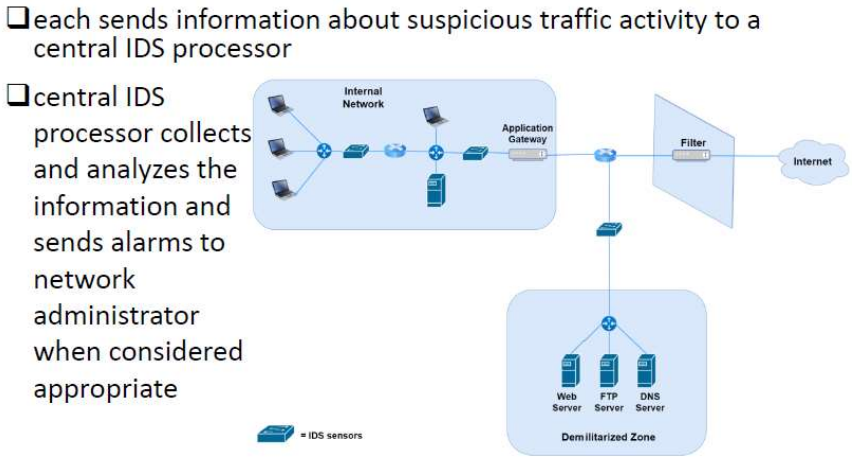

So, this is the high-security zone. So, there is a packet filter at the boundary of the network, and then there are several routers, desktop computers, servers, and so on in the network. And there are three IDSs: one, two, and three. So, these are three IDS sensors which analyze the packet headers as well as application data of the packets that flow through these IDS sensors. When multiple sensors are deployed, each sends information about suspicious traffic activity to a central IDS processor.

So, for example, this IDS sensor analyzes the packets that are flowing to these three devices. So, it analyzes these packets, which flow to these three devices, and if it finds some suspicious activity in the packets, then it sends information about such activity to a central IDS processor. This IDS sensor analyzes the traffic that flows to and from all these servers, and it sends information about suspicious activity to a central IDS processor, and so on and so forth. The central IDS processor collects and analyzes the information from all the IDS sensors and sends alarms to the network administrator when it is considered appropriate. So, the IDS processor analyzes the information that it gets from all the IDS sensors spread throughout the network.

So, that provides an exhaustive picture of the activity that's happening in the network. If it detects some suspicious activity, then it sends an alarm to the network administrator. And the network administrator can then take appropriate actions. So, in this figure, as we said earlier, the network is partitioned into two regions. One is a high-security region; that is this one.

So, this high-security region is protected by a packet filter, which is this one. And application gateway, which is shown here. And monitored by two IDS sensors, which are this one and this one. Then the other part of the network is a low-security region called demilitarized zone, or DMZ. That is shown here.

So, this is the demilitarized zone. It is protected by the same packet filter and monitored by one IDS sensor. So, this is the IDS sensor, which monitors the traffic to and from the demilitarized zone. The DMZ contains the organization's servers that need to communicate with external users. So, this web server may, for example, host the company's website, and external users can connect to this web server and access the company's webpage.

Similarly, this FTP server might host the files that belong to this company, which are non-confidential and which are publicly accessible. And external users may connect to this FTP server and access those files. So, this demilitarized zone has servers which provide data that can be accessed by external users. A node in the external network can only access nodes in the DMZ. So, an external user on the internet can only access nodes in the demilitarized zone.

It cannot access any of the nodes in this high-security part of the organization's network. The firewall blocks all access to the high-security region. So, for example, if a user tries to send some packets to one of the nodes in this high-security region, then this packet filter

will block those packets. But if the user tries to send packets to the demilitarized zone, then the filter will allow those packets to pass. So, in this picture, this is the demilitarized zone.

A demilitarized zone contains and exposes an organization's external phasing services, for example, web server to the public internet. It is isolated from the rest of the internal network using IDSs. So, for example, we can see that this DMZ is isolated from this high-security region via this IDS sensor and this IDS sensor. So, what is the reason for such isolation? The reason is that since machines in the demilitarized zone are accessible to the public, they are the most likely machines to be compromised in the entire network.

So, packets from the internet flow into the demilitarized zone, and some of these may be malicious packets. They may infect servers in the demilitarized zone. Hence, machines in the demilitarized zone are most likely to be compromised. So, machines in the high-security region can also be compromised, but that can happen only if there is some malicious employee in the company or there is some flash drive which is used in the internal network and that is some malicious information. But these events are rare.

The most likely way in which a host can get infected is if some users in the internet send some malicious traffic. But such malicious traffic can only reach the demilitarized zone, not the high-security region. Thus, machines in the DMZ are accessible to the public, and hence they are the most likely machines to be compromised in the entire network. The ideas can protect machines in the rest of the internal network from being compromised if a machine in the demilitarized zone is compromised. So, suppose this web server gets compromised; then the malware that has infected this web server may try to spread itself by sending information to hosts in this high-security region, but then those packets will be blocked by this IDS sensor and/or this IDS sensor.

So, the infection does not spread from the demilitarized zone to the high security region, because of the presence of these IDS sensors, this one and this one. So, that's a reason for having IDS sensors at the boundary between the demilitarized zone and the high security region. So, in this picture, we know that there are three IDS sensors: one, two, and three. So, in this example organization's network, there are three IDS sensors. Why not use only one IDS sensor, which could be placed just behind the packet filter or combined with it?

So, we could have only one IDS sensor, which is either combined with this filter or connected to it at this point. So, why do we need multiple IDS sensors in a network? So, the reason is the following. How does an IDS function? Typically an IDS compares each passing packet with tens of thousands of data patterns known as signatures.

For example, a signature may be a part of a virus or a worm. So, if the information in the packet matches that signature, then it indicates that it is malicious. So, that's how an IDS functions. It compares each passing packet with tens of thousands of signatures which are stored in the IDS. And if it detects some suspicious activity, then it sends an alert or blocks the packet.

So, this comparison of the packet information with signatures requires a significant amount of processing, especially if the organization's network receives a large amount of traffic from the internet. So, if it's a busy network, it will receive a lot of traffic from the public internet. In that case, all these packets will have to be compared with signatures—tens of thousands of signatures, which requires a lot of processing. So, by placing the IDS sensors further downstream, each sensor only sees a fraction of the organization's traffic and can more easily keep up. So, if the IDS sensor were kept here, then all the traffic flowing into the organization's network would pass through the IDS.

So, a large amount of traffic would need to be analyzed by the IDS sensor. But if it is kept downstream, for example, this is a downstream point which receives only a fraction of the traffic that comes into the organization's network. So, this is a downstream point, and this is another downstream point. So, by keeping these IDS sensors downstream, they receive only a fraction of the traffic that flows to the organization's network, and hence it is easier for them to cope with the traffic, analyze it in real time, and raise alerts. So, by placing the IDS sensors downstream, each sensor only sees a fraction of the organization's traffic and can more easily keep up with the traffic.

So, that's the reason for having multiple IDS sensors. And each of the IDS sensors analyzes only a part of the traffic that flows into the organization's network. Hence, we typically use multiple IDS sensors. We don't use just one IDS sensor combined with the packet filter or connected to it. So, in summary, we introduced intrusion detection systems.

We discussed that an intrusion detection system may either block traffic or allow traffic to pass. And we looked at the structure of an example network in which IDS sensors are used. And we introduced the concept of a demilitarized zone. We will continue our discussion of IDSs in the next lecture. Thank you.