Network Security

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

Week - 09

Lecture - 51

Firewalls and Intrusion Detection Systems: Part 1

Hello, in this lecture and the next several lectures, we will discuss firewalls and intrusion detection systems. These are devices which defend the network of an organization against attacks by malicious users. We know that most organizations, such as universities and companies, have their networks connected to the public internet. These organizations have a local area network, which is used by all the employees or students and so on of the organization. But these local area networks are connected to the public internet.

So, taking advantage of this connection to the public internet, attackers who are connected to the public internet may attempt to attack the organization's network in various ways. For example, they may infect machines on the organization's network with malware such as viruses, worms, Trojans, and so on. They may try to steal the secrets of the corporate. This will happen if some confidential information leaks out from the organization's network and is accessed by attackers, or attackers may also attempt to map the internal network configurations. That is, they attempt to figure out how the different nodes in the organization's network are connected to each other.

And then they can later on take advantage of this knowledge to attack the organization's network. By mapping the internal network configurations, they know how different nodes are connected to each other, what the topology of the network is, and so on. So, they can then make use of this knowledge to attack the organization's network. Attackers can also launch denial of service attacks on the organization's network. This can be done after mapping the internal network configuration.

In a denial of service attack, a large number of bogus packets may, for example, be sent into the network. So, as a defense against these kind of attacks, we'll discuss firewalls and intrusion detection systems. These are devices which can be used to detect and/or prevent attacks which are of the kind that are listed here. This shows an example of a firewall. This is an organization's network.

It consists of several routers such as the ones shown here, and several computers such as the one shown here, several servers, and so on. These are all connected to each other. There may be hundreds or thousands of such nodes that are connected to each other in the organization's network. And a firewall is a device which typically sits at the boundary between the organization's network and the public internet. And all the traffic that flows from the organization's network to the public internet and from the public internet to the organization's network, all this traffic flows through the firewall, and it is monitored by the firewall.

And it blocks certain traffic and only allows the rest of the traffic to flow through. Depending on what policies are configured by the system administrator, certain packets are blocked and only other packets are allowed to pass through. For example, if the firewall suspects that certain packets that are being sent from the public internet to the organization's network, they are malicious packets, then the firewall may block those packets. And conversely, if the firewall suspects that some packets are being used to leak confidential information from the organization's network to the public internet, then the firewall may block those outgoing packets as well. Depending on what security policy is configured by the system administrator, the firewall blocks certain packets to defend against such attacks.

First, we'll discuss firewalls in detail, and then we'll discuss intrusion detection systems. Intrusion detection systems are also similar to firewalls, but there are some differences, as we'll see. So, intrusion detection systems are similar to firewalls in the sense that they also monitor the traffic that is flowing through the network, and they block certain packets and they allow the remaining packets to flow through. But there are some differences, which we'll discuss later. A firewall, which is shown over here, is a combination of hardware and software.

It controls access between an organization's internal network, which is shown over here on the left side, and the internet, which is shown on the right. It allows some packets to pass and blocks the other packets based on a given security policy. As an example, there may be a particular website in the organization's network which should be only accessed by the employees of the organization's network. If a user from the public internet tries to access that private website, in that case the firewall may block that request for a connection to the private website. This is an example where the firewall blocks certain packets.

As another example, suppose there is a public website in the organization's network which is allowed to be accessed by external users. In that case, the firewall allows packets which request access to the public website. So, a firewall allows some packets to pass and blocks the other packets based on the given security policy. By this process of blocking some packets, the firewall prevents intruders from attacking the internal network. It also prevents confidential internal data from leaving the organization's network.

This is done by blocking certain packets which are flowing out from the organization's network to the public internet. We now discuss the properties of a firewall. All traffic from outside to inside, and vice versa, passes through the firewall. That is, all traffic from the public internet to the organization's network and all traffic from the organization's network to the public internet. All of this traffic passes through the firewall and is monitored by it.

Only authorized traffic is allowed to pass. What is meant by authorized traffic? Whatever the security policy configured by the network administrator considers as being okay to be passed through the firewall, that is authorized traffic. So, only authorized traffic is allowed to pass, and unauthorized traffic is blocked by the firewall. Clearly the firewall should not be compromised because if the firewall is compromised by a malicious user, then it will allow malicious traffic to pass, and it may block certain legitimate traffic.

So, the firewall itself is designed and maintained such that it is difficult to compromise. For example, unnecessary services or applications on the machine are removed because these unnecessary services may have some bugs or malware in them. So, it's better to remove them. And newly available security patches are installed expeditiously. There are often new kinds of malware that are released by attackers.

So, to defend against these malware, we require security patches. Newly available security patches are installed expeditiously in firewalls to make it difficult to compromise firewalls. A firewall may be implemented in hardware as a standalone device, or it may be in software on a PC. But clearly, if it is implemented in hardware as a standalone device, then it will perform better than a firewall which is installed in software on a PC. So, firewall which is installed in hardware as a standalone device will have a better performance and it will be able to process more packets per unit time than a firewall which is installed in software on a PC.

On the other end, a firewall which is installed in software on a PC will be less expensive than a firewall which is implemented in hardware as a standalone device. Also, many routers support basic firewall functionality. So, an organization can purchase a router and

get some basic firewall functionality with it. So, we now provide some background about TCP/IP protocols, which is required to better understand firewalls and intrusion detection systems. So, we'll discuss a variety of protocols such as ICMP, DNS, and so on.

So, later we'll see how the headers pertaining to these protocols are examined by firewalls, and this knowledge is used to block certain traffic while allowing other traffic to pass. So, we start with the Internet Control Message Protocol (ICMP). ICMP is a protocol used by hosts and routers to communicate network layer information to each other. So, this is a layer three protocol, that is the network layer; it's a network layer protocol. A typical use of ICMP is for error reporting.

So, ICMP does not deal with data packets, but as the name suggests, it deals with control messages, such as reporting errors and so on. One example is this. Suppose a source host A wants to send a packet to a destination host B and while forwarding the packet, if an IP router is unable to find a path to the destination address B, then it sends an ICMP packet to the source that is A, indicating the error. So, this is one example. Recall that there is a routing table in each router, which provides the next hop router for each destination address.

But it can happen that a particular IP router is unable to find a path to the given destination address. Then it sends an ICMP packet to the source, indicating the error. So, this way, the source knows that the packet could not reach the destination. So, what does the source do? The source may display a message such as "Destination host unreachable" or "Destination network unreachable" to the end user.

- A typical use of ICMP is for error reporting
  - e.g., while forwarding a packet, if an IP router is unable to find a path to the destination address, then it sends an ICMP packet to source indicating the error
  - may result in display of "Destination host unreachable" or "Destination network unreachable" message to end user
- ICMP packets have the "Protocol" field in the IP header equal to 1
  - note: this field equals "6" for TCP packets and "17" for UDP packets
- Examples of ICMP packets:
  - a "*redirect*" packet, which tells source host to use a particular router for forwarding to a particular destination, presumably because the router the source chose on a previous packet was not the best path to the destination
  - a "*ping*" packet, which is supposed to be echoed back by the system that receives it

So, the end user knows that there was some problem in reaching the destination address. So, in this way, ICMP provides useful information to the source host in this case. So, indicating that the destination was unreachable. ICMP packets have the "Protocol" field in

the IP header set to 1. So, the "Protocol" field indicates what kind of payload is carried by the IP packet.

So, this field, that is, the "Protocol" field, equals 6 for TCP packets and 17 for UDP packets. So, by examining this field, one can find out whether the packet carries ICMP information, or TCP information, or UDP information, and so on. Some examples of ICMP packets are as follows. One example is a redirect packet. So, a redirect packet tells a source host to use a particular router for forwarding to a particular destination.

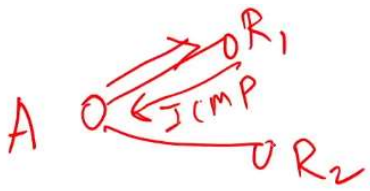

That typically happens because the router the source chose for a previous packet was not the best path to the destination. So, let's illustrate this with a simple diagram. So, this is a source A, which sent a packet initially to router R_1. And there is another router, R_2. This source A wants to reach a particular destination B. It initially sends a packet to router R_1.

But then router R_1 finds out that to reach B, source A should actually send the packet to R_2, not to R_1. So, R_1 sends an ICMP packet to node A. This ICMP packet indicates to node A that to reach node B, the best router is router R_2, and it is not router R_1. Then, subsequently, node A sends the same packet through router R_2 instead of R_1. So, this is one example of an ICMP packet. In this case, the ICMP packet is a "redirect" packet.

It tells the source host A to use a particular router, R_2, for forwarding to a particular destination, that is, B, because the router that the source chose on a previous packet, namely R1, was not the best path to the destination. Another example of an ICMP packet is a "ping" packet. We often use "ping" packets in debugging. A "ping" packet is sent by one host to another host, and it is supposed to be echoed back by the system that receives the "ping" packet. So, this can be used by the source host to check whether the destination host is alive and reachable.

So, it is used for seeing if a system is alive and reachable. So, if a system can be pinged, then it is alive and reachable, and vice versa. So, now we discuss some examples of attacks using ICMP messages, and later on we'll discuss how firewalls can be used to defend

against such attacks. So, how can an ICMP "ping" packet be exploited by an attacker? It can be exploited by an attacker to find machines to break into.

So, the attacker can send "ping" packets to different possible IP addresses, and whenever it gets a response, it means that there is a machine with that IP address. So, the attacker can later on try to send some malicious packets to that machine. So, "ping" can be used by an attacker to discover machines which can be broken into. Another example is an ICMP message can be sent to an internal host falsely claiming that some range of addresses is unreachable will cause Alice to end its connections to machines in the range specified by that ICMP message. So, we discussed that one of the types of ICMP messages is a destination unreachable message.

So, it informs the source that a particular destination address or range of addresses is unreachable. But this can be maliciously used. An intruder can send an ICMP message to an internal host, say Alice, claiming that some range of addresses is unreachable. So, this is a fraudulent packet. In fact, the range of addresses is actually reachable, but the attacker falsely claims that this range of addresses is unreachable, and this causes Alice to end its connections to machines in the range specified by that ICMP message.

So, this is a form of denial of service, so Alice is not able to connect to machines in that range of IP addresses. Another example of attacks using ICMP messages is, ICMP redirects can be used to cause a host to send traffic in a different direction, possibly toward a compromised machine. So, we discussed an example earlier of the redirect message. So, consider this node A, and this is R_1 and this is R_2. In this case, suppose A sends a packet initially to R_1.

So, R_1 has been compromised by an intruder who controls router R_2. So, R_1 sends an ICMP packet that is a redirect packet to A, which causes A to send its packets to router R_2, which is controlled by an intruder. So, subsequently, A sends all its packets to router R_2, which can then launch some malicious attacks on these packets that are sent by A to R_2. So, this is an instance where ICMP redirects are used to cause a host to send traffic in a different direction, possibly toward a compromised machine. So, this allows man-in-the-middle attacks to take place.

So, in this example, possibly A wants to communicate with some other hosts, say B, but this redirect-based attack causes all the traffic between A and B to flow through the router R_2. So, R_2 can act as a man-in-the-middle and launch attacks on the traffic that flows between A and B. This is possible because of an ICMP redirect packet. So, now we will

discuss another program called Traceroute. Traceroute is a program that can be used to trace a route from a host to any other host in the world. For example, consider this host A, which communicates with another host B through several intermediate routers.

So, traceroute allows A to find out the IP addresses of all the routers on the path from A to B. So, A can get the IP address of this router, this router, this router, and so on. So, it can get the IP addresses of all the routers on the path from A to B using the traceroute program. So, how does traceroute work? We'll discuss this. So, traceroute provides the IP addresses of all the routers on the path from a given source to a given destination.

It is implemented using ICMP messages. So, when the traceroute program is used, the source A sends a series of ordinary UDP packets to the destination B. But these are special packets. They are special because each packet contains an unlikely UDP port number; that is, the destination port number in these UDP packets is the port number of some port at which no process is running. So, this is the destination port number is one at which there is no corresponding application process. Also, this series of ordinary UDP packets has the feature that the first of these UDP packets has a TTL of 1, the second has a TTL of 2, the third has a TTL of 3, and so on and so forth.

So, TTL is time to live. So, whenever a source node sends a packet into the network, it initializes a field called TTL or Time To Live in the packet. When a source node sends an ordinary data packet, then typically the TTL is initialized to some value like 64 or so. So, what is the purpose of this TTL packet? So, each time the TTL packet reaches a router, the router decrements the value of TTL.

So, for example, if A sends a packet to the next router, that router will decrement TTL. It will reduce from 64 to 63. Then, when that router forwards it to the next router, it will decrement it from 63 to 62, and the next router from 62 to 61, and so on. Each time the packet reaches a router, the TTL value is reduced by 1. And when the TTL value reaches 0, the router drops the packet. So, the purpose of the TTL field is that it helps in preventing packets from indefinitely revolving in routing loops. So, there can occasionally be routing loops in the network. So, this is an example of a routing loop.

So, if a packet reaches one of the routers in this loop, in that case, suppose the packet reaches this router; then the destination is somewhere here, but this packet keeps on revolving around this routing loop because in the routing tables of these routers, the next hop for this router is this router; the next hop for this router is this router; the next hop for this router is this one, and so on and so forth. So, the next stop for this router is this router.

So, the packet keeps on revolving around this loop instead of reaching its destination. So, to prevent a packet from indefinitely revolving around such routing loops, the TTL packet is used. It is initialized to some value like 64, and it is decremented by 1 each time a packet is forwarded by a router.

So, when it reaches 0, the packet is dropped. So, if a packet gets stuck in a routing loop, then eventually it will be dropped because of the TTL field reaching 0. So, that's the purpose. So, there can be such routing loops temporarily in a network because of the use of distributed routing algorithms such as RIP. So, to prevent a packet from indefinitely revolving around such routing loops, we use the TTL packet.

Fine, so now this TTL packet is used, and traceroute uses this TTL packet to trace a route from the source to the destination. So, in this series of ordinary UDP packets that is sent by the source host to the destination host, the first UDP packet has a TTL of 1, second has a TTL of 2, third has a TTL of 3, and so on and so forth. When the n'th packet arrives at the n'th router, then the n'th router observes that the TTL has just expired; that is, the TTL has reached 0. So, we call that the n'th packet has a TTL of n. So, when the n'th packet reaches the n'th router, the TTL will just become 0. So, the n'th router will observe that the TTL has just expired.

- The first of these UDP packets has a TTL of 1, the second of 2, the third of 3, and so on
- When the n'th packet arrives at the n'th router:
  - the n'th router observes that the TTL has just expired
  - according to the rules of the IP protocol, the router discards the packet and sends an ICMP message to the source
  - this message includes the IP address of the router
- When this ICMP message arrives at the source host, it obtains the IP address of the n'th router on the path to the destination host

So, according to the rules of the IP protocol, the router will discard the packet and it sends an ICMP message to the source. So, this ICMP message informs the source that the packet that the source had sent has an expired TTL. The TTL reached 0; so, hence, the router has discarded the packet. So, this information is provided to the source host by this ICMP message. This message includes the IP address of the router which discarded the packet.

So, when this ICMP message arrives at the source host, it obtains the IP address of the n'th router on the path to the destination host. So, we can see in this example that for this router, n =1; for this router, n =2; for this one, n =3; for this one, n =4; and so on. So, this is the router with n =1, this one with 2, 3, 4, and so on. So, when the first packet reaches the first router, it drops the packet and sends an ICMP message to A. At that point, A comes to

know the IP address of this router, 1. Because from the source IP address of the ICMP message, A gets to know the IP address of this node one.

When the second packet reaches router 2, its TTL expires, and router 2 sends an ICMP message to source A. At that point, A comes to know the IP address of router 2, and so on and so forth. When the third packet reaches router 3, router 3 decrements the TTL, and it becomes 0. It sends an ICMP message to source A, and then A comes to know the IP address of router 3. So, in this way, the source host A comes to know the IP address of each router on the path from itself to the destination B. So, this process continues until one of the UDP packets sent by the source host reaches the destination host, that is B in this example.

So, recall that each of these packets contains an unlikely UDP port number. When one of the UDP packets sent by the source host reaches the destination host B, B will find out that there is no application process listening at the destination port number. So, it will send an ICMP packet to the source host A, informing that the destination port number is unreachable. So, from this, A will get to know the IP address of node B. So, in this way, it has obtained the IP addresses of all the routers on the path from A to B. So, it already knew the IP address of node B, but it didn't know the IP addresses of the intermediate routers.

So, those it gets to know through the traceroute program used in this way. So, this is the legitimate use of the traceroute program, but the traceroute program can be used by an attacker to attack an organization's network in the following way. It can map the internal configuration of the organization's network. So, it can keep on sending such special UDP packets to nodes inside the network. And it gets to know how different nodes are connected to each other.

So, it can map the internal configuration of the organization's network in this way. And it can use this knowledge to later attack the organization's network. So, it can use the configuration that is obtained to later attack the organization's network. So, in this way, the traceroute program is actually useful. It can be used to find the IP addresses of the routers on the path from the source to a destination.

But it can be used by an attacker to attack an organization's network. Later on, we'll discuss how firewalls can be used to defend against such attacks based on the traceroute program. So, in summary, we started our discussion of firewalls and intrusion detection systems. We first discussed firewalls, and later we'll discuss intrusion detection systems. So, to

understand firewalls better, we need to understand some protocols of the TCP/IP protocol stack in detail.

So, we discussed various protocols such as ICMP, and then we discussed the traceroute program, which is based on ICMP packets. We will continue our discussion of firewalls in the next lecture. Thank you.