Network Security

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

Week - 06

Lecture - 33

IPsec and Virtual Private Networks (VPNs) for Network-Layer Security: Part 2

Hello, in the previous lecture, we started our discussion of IPsec and virtual private networks. We now continue the discussion. So, we discussed last time that security associations have been created between R1 and R2, which are routers that want to communicate with each other. So, now the question is, how is data transferred over these security associations? So, data is transferred over SAs using a protocol known as Encapsulating Security Payload (ESP).

So, ESP is an IPsec protocol that provides confidentiality, message integrity, and detection of deleted, duplicate, or reordered packets that are sent over a security association. So, these mechanisms are provided by the usual means. Confidentiality is provided using encryption. For achieving message integrity, a message authentication code (MAC) is added to each packet. So, for a message m, the message authentication code is H(m,s), where H is a cryptographic hash function and s is the authentication key.

And a sequence number is used to detect addition, deletion, or reordering of packets by intruders. So, for every packet that is sent over an SA, a sequence number is added, and the sequence number typically increments for every packet, so the sequence number is an increasing sequence, and this can be used to detect addition, deletion, or reordering of packets. So, we have discussed the use of sequence numbers for these attacks: addition, deletion, and reordering of packets, so in the usual way, these sequence numbers are used to defend against these attacks. ESP adds a header which contains the SPI, that is, the security parameter index, which we discussed in the previous lecture. And the header also contains the sequence number, message authentication code, and some other fields.

So, there is the original packet which is to be transferred, and ESP adds a header which includes all this information to the packet. So, IPsec has the following modes. One is the transport mode, and the other is the tunnel mode. And each security association operates in one of these modes. So, the tunnel mode is used for creating virtual private networks.

And the transport mode is very similar to the use of SSL. We first briefly discuss the transport mode. In the transport mode, ESP's payload data is a message for a higher layer protocol, such as UDP or TCP. So, ESP is at the network layer, and the layer above it is the transport layer. So, in the transport mode, ESP's payload data is a message for a higher layer protocol, such as UDP or TCP.

And in this case, IPsec acts as an intermediate protocol sub-layer between the transport layer and the network layer, just as SSL acts as a sub-layer between the transport layer and the application layer. So, IPsec is just an intermediate protocol sub-layer, which transfers the packets that are sent by the layer above it, that is, the transport layer. When an ESP packet is received, its payload is passed to the higher layer, that is, the transport layer in this case. So, this is the transport mode. It is very similar to the use of SSL.

So, SSL transports the application data, which is provided by the application layer, securely. Similarly, the transport mode transfers the transport layer data, provided by the transport layer, securely. So, this is the transport mode. The other mode of IPsec is tunnel mode. It is commonly used for creating virtual private networks.

So, this shows the format of an ESP packet when the tunnel mode is used. In this case, the ESP's payload data is itself a complete IP packet. So, this shows an original IP packet, which is to be securely transferred over a security association. So, this original IP packet has an IP header, TCP or UDP header, and the payload data, that is, the application layer data. So, this is the original IP packet, which is to be transferred.

- When an ESP tunnel mode packet is received by a node, its payload is forwarded on as a normal IP packet

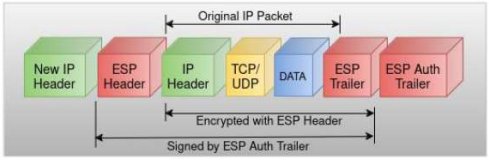

Now, this is enclosed in another IP packet. So, specifically, one header known as the ESP header is added, which contains some information, which we'll discuss later. And there is a trailer, which contains some additional information. And this is the message authentication code, shown here as 'ESP Auth Trailer'. So, this is the message authentication code.

And all these fields, that is, the original IP packet and the ESP trailer, are encrypted using information that is there in the ESP header. All this information, including the ESP header and the original IP packet and the ESP trailer; all this is signed by the ESP message integrity algorithm and the MAC is put in this place. So, this is the MAC that is provided. And a new IP header is added to this, to these fields. So, this new IP header has its own source IP address, destination IP address, and other fields in the IP header.

When an ESP tunnel mode packet is received by a node, its payload is forwarded on as a normal IP packet. So, in the example that we discussed in the previous lecture, where there were tunnels created between the headquarters gateway router and the branch office gateway router. So, in this case, there is one SA from the headquarters gateway router to the branch office gateway router and one SA from the branch office gateway router to the headquarters gateway router. So, when one of these routers receives an ESP tunnel mode packet from the other end, it extracts the inner IP packet, that is this one, and then forwards it as a normal IP packet into the network. So, for example, if the branch office gateway router receives an ESP tunnel mode packet from the headquarters gateway router, then it will extract the inner IP packet.

So, it extracts the inner IP packet and forwards it into its local area network. Here's an example. This is the headquarters office, and this is the branch office, and there is a security association from R1 to R2 as shown here. So, this is a part of a virtual private network. Now, suppose router R1 receives a packet from host 172.16.1.17 in the headquarters network.

- Suppose router R1 receives a packet from host 172.16.1.17 (in headquarters network) destined to host 172.16.2.48 (in branch-office network)
- At router R1, this packet becomes the payload of an ESP message sent over the SA to R2
- Router R2 unwraps the payload IP packet and forwards it to 172.16.2.48

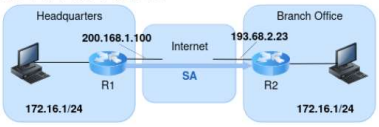

So, this source host is somewhere here in this headquarters network. And the destination is host 172.16.2.48, which is somewhere in the branch office network. So, that is somewhere over here in this network. Then, at router R1, this packet becomes the payload

of an ESP packet that is sent over the security association to R2. And then, router R2 unwraps the payload IP packet and forwards it to the destination, that is, 172.16.2.48.

So, in this example, notice that there is a virtual private network created, which includes all the machines in the headquarters office and all the machines in the branch office. So, all these machines in the headquarters office and in the branch office, they all have IP addresses starting from this prefix, 172.16. some other numbers. So, the prefix is common for all the machines in the headquarters office as well as the branch office. So, we can see that the prefix is 172.16.1.1/24 and the prefix is 172.16.1/24 in each case. So, the prefix here is 172.16.2/24, whereas here it is 172.16.1/24. But the common prefix is the same, that is, 172.16.

But on the internet, other IP addresses with other prefixes are used, that is, in the public internet, the IP address of this interface is, for example, 200.168.1.100, and the IP address of this interface is 193.68.2.23. So, in this example, where a source node somewhere in the headquarters office sends a packet to a node in the branch office network. So, the source IP address is this, 172.16.1.17, and the destination IP address is this, 172.16.2.48. So, in the packet created by that source node, which is in the headquarters office, these are the source and destination IP addresses. Now, this packet is encapsulated in another IP packet, so that was shown in the previous figure.

So, the packet that was created by the headquarters source node, that is this original IP packet and that is encapsulated inside this outer IP packet and there is a new IP header that is added in this outer IP packet and in that IP header, the source and destination addresses are these; 200.168.1.100 and 193.68.2.23. So, these are the IP addresses in this outer IP header. This complete packet is tunneled over the public internet, and then at the destination, namely at the branch office gateway router, this original IP packet is extracted from this outer IP packet. Then, this original IP packet is forwarded as a usual IP packet inside this branch office network. So, that is how this tunneling is done. To summarize this tunneling, the source node, which has the IP address 172.16.1.17, sends a packet to this gateway router.

Then, this gateway router puts that packet inside another IP packet and sends it over the public internet to this R2. This R2 extracts the inner packet and then sends that packet to this destination, which is in the branch office network. So, this is known as tunneling. So, the source as well as the destination use the IP address numbering of the local area network

of the company. That is 172.16.something. Whereas, in the public internet, some other IP addresses are used, which are shown here.

The examples are shown over here. Now, we discuss the format of an ESP packet. This figure shows an ESP packet when the tunnel mode is used. So, in this case, this is the original IP packet. So, the original IP header is here, and the original IP datagram payload is over here.

And a trailer is added to it, which contains all this information. There is padding added, and then the length of the pad and the next header. For example, the next header might be a transport layer header, either TCP or UDP. So, all this information is added in a trailer. So, this original IP packet and the trailer are encrypted using the algorithm and key corresponding to the security association.

- ESP header containing SPI and sequence no. added
- A MAC, computed over (ESP header + encrypted information) using the algorithm and key corresponding to SA, added
- Ordinary IP header added
- Recall: sequence numbers not included in packets in SSL; why is sequence number included in ESP?

❑ sequence numbers were included in TCP; hence they could be omitted from SSL

❑ sequence numbers not included by IP; hence, included by ESP

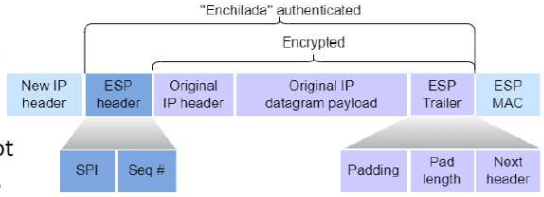

So, that is shown over here. This information is encrypted using the algorithm and key corresponding to the security association. And then an ESP header, which contains the security parameter index and sequence number, is added. So, that is shown over here. This is the ESP header, which contains the security parameter index and sequence number.

So, all these are headers; I mean, this is additional information added, which includes SPI, sequence number, padding, pad length, and next header. But some of this information is encrypted, and the other is not encrypted. So, this information is encrypted, and this information is not encrypted. So, the reason that this information is not encrypted, for example, the SPI is not encrypted, is because the receiving router needs to know what the SPI is to find out how to decrypt this packet. So, this should not be encrypted.

So, for this reason, the SPI is not encrypted. Now, a MAC that is computed over the ESP header and the encrypted information is calculated using the algorithm and the authentication key corresponding to the security association. So, this MAC is computed and added. So, this is the MAC that is added to the packet. This is the ESP MAC.

And then an ordinary IP header is added, which is this IP header. So, this has the source IP address and destination IP address of the public internet. Now, one observation is that sequence numbers are not included in packets in SSL. So, why are sequence numbers included in ESP? So, we can see that the sequence number is included in the ESP header, whereas sequence numbers were omitted from SSL packets.

Instead, the sender and receiver kept track of the sequence numbers using counters. So, why are sequence numbers included in ESP, whereas they were not included in SSL? So, the reason is that the sequence numbers were included in TCP, over which SSL runs. Hence, sequence numbers could be omitted from SSL. And in contrast, ESP packets are sent over IP, but IP does not have any sequence numbers.

Hence, in order for the receiver to put the packets in the correct order, ESP adds sequence numbers to enable the receiver to arrange the packets correctly. So, that's the reason for including sequence numbers in ESP but not in SSL. And one thing to note is that the sequence number is not encrypted, so it is sent in the clear. So, what is the reason for not encrypting the sequence number? So, the reason is that an intruder, who is eavesdropping on all the packets sent between R1 and R2, that intruder can keep track of sequence numbers using a counter.

So, the intruder can predict what the sequence number in each packet should be. This is ignoring any packet losses, duplicate packet transmissions, and so on. So, the intruder can mostly predict what the sequence number should be. Hence, the intruder knows the plaintext, which is the sequence number, and if the sequence number were encrypted, the intruder would get the corresponding ciphertext since the sequence number is encrypted. Hence, the intruder would get a plaintext and a corresponding ciphertext, and it becomes easy to break a cipher if the intruder gains access to a lot of plaintext and corresponding ciphertext.

So, for this reason, the sequence number is not encrypted. So, we see that for this reason, the sequence number is put in the header, which is not encrypted. And similarly, the SPI is required for the receiver to interpret how the packet is to be decrypted and how the MAC

is to be verified. So, for this reason, the SPI is also not encrypted. So, this concludes our discussion of ESP. We now discuss virtual private networks.

So, first, we will discuss private networks, and then we'll discuss virtual private networks. So, the context is that many companies have offices in multiple locations. For example, a company might have an office in Mumbai, one in Bangalore, one in Hyderabad, and so on. And these offices are often in different countries or different towns. So, before the public internet appeared, such companies used to lease communication lines from telephone companies to connect their offices at different locations.

So, there were dedicated communication lines between the offices at different locations. And some companies still do this. They connect their offices at different locations using communication lines leased from telecom companies. Such a network is called a private network. So, it's owned by the company.

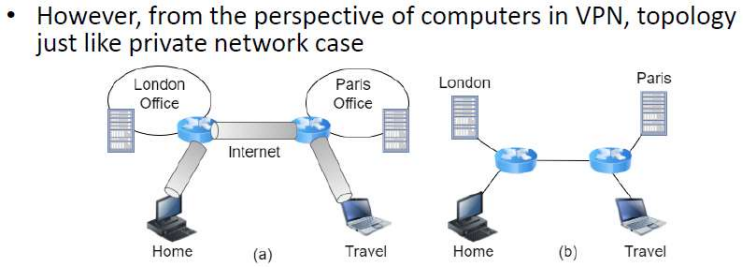

So, it's a private network. So, the traffic of other users, apart from the company employees, that other traffic cannot flow over this private network. So, as we can easily imagine, private networks are very secure. Intruders need to physically tap lines to obtain confidential information, which is difficult. So, in the case of the internet, the routers and communication links are shared among many users, so the traffic of many users flows over the public internet.

But in the case of private networks, only the traffic of the company employees flows over private networks. Hence, private networks are very secure to obtain confidential information intruders need to physically tap lines, which is very difficult. But there is a shortcoming of private networks that the shortcoming is cost leasing dedicated lines between two points is expensive, so for this reason, only very big companies can afford to set up private networks. Hence, when the public internet appeared, companies wanted to use it to connect offices at different locations, but with strong security, just like in a private network. So, the objective was that communication would occur over the public internet,

but from the security point, it should be as secure as a private network. So, this would reduce the cost as well and still provide strong security.

So, that led to the invention of virtual private networks. So, this illustrates a virtual private network. So, a common VPN design is shown in these figures. So, in this example, a company has an office in London and an office in Paris. Each office is equipped with a gateway router.

This is the gateway router of the London office, and this is the gateway router of the Paris office. Tunnels are created through the internet between each pair of gateway routers, just like this tunnel over here. So, there is a tunnel between each pair of gateway routers, and tunnels are also created between gateway routers and employees who are either traveling or working from home. So, for example, there is a tunnel between this gateway router and this employee who is working from home, and there is a tunnel from this gateway router to this employee who is traveling. And how is a tunnel created?

It is created by establishing two IPsec security associations, one in each direction. So, to create this tunnel, for example, there is one IPsec security association from this gateway router to this gateway router and one security association from this gateway router to this gateway router. So that's how a tunnel is created. Now, a VPN is more flexible than a private network that is built using leased lines. The reason is that the tunnels can be set up on demand to any employee who is traveling or working from home with an internet connection.

So, employees may work from hotels, for example, when they are traveling. So, in this case, a tunnel needs to be set up to the hotel computer. So, using a private network, we cannot connect users securely and make them part of the private network. So, a private network is static. It is between different offices of a company.

So, a private network cannot connect an employee who is traveling and is connecting from some hotel to the private network. So, these are limitation of private networks but VPNs are more flexible because it can create tunnels to users who are connecting from home or who are traveling as well. From the perspective of computers in the VPN, the topology is just like a private network case. So, this is the topology from the perspective of users in the London office, in the Paris office, and from the perspective of the employees who are at home or who are traveling. So, it appears as though they are all part of the same local area network.

So, this view is achieved by tunneling, which we discussed earlier in the case of ESP. So, recall that these nodes, which are in the London office and which are in the Paris office, and these employee computers, they just send normal IP packets using the source IP and destination IP addresses of this LAN, of the company LAN. But then, these IP packets are encapsulated in outer IP packets and are sent over the public internet, where the inner IP packet is extracted and then sent to the destination. So, using this tunneling process, from the perspective of computers in the virtual private network, the topology is just like the private network case. But actually, the connectivity is over the public internet.

So, what security properties does a VPN provide? So, a VPN provides encryption and message integrity because IPsec is used. So, we have discussed that IPsec provides encryption and message integrity; hence, a virtual private network provides these mechanisms. Also, all traffic between a given pair of gateway routers can be aggregated into two security associations, one in each direction. So, for example, there might be hundreds of computers in the London office and hundreds of computers in the Paris office.

The traffic that is flowing between these computers, all that traffic is aggregated into only two security associations, one in each direction. One from this gateway router to this one and one from this gateway router to this one. So, an advantage of doing this is that intruders on the public internet cannot find out the amount of traffic that flows between any pair of machines. So, there might be hundreds of computers in this office. So, all that traffic is aggregated and then just sent on one connection from, that connection is the security association.

So, all that traffic is aggregated and that is sent over only one connection from this gateway router to this gateway router, and hence there's only one flow from this gateway router to this one and one flow from this gateway router to this one. Any interceptors who are sniffing that traffic cannot find out how much traffic is flowing between a given pair of machines, one in this office and one in this office. So, they can only observe the aggregate traffic that is flowing between these gateway routers. So, they cannot individually find out how much traffic is flowing between a pair of machines in different offices. So, that's the advantage of virtual private networks.

So, hence, a virtual private network defends against traffic analysis attacks. So, in our lecture on different attacks on networks, we discussed traffic analysis attacks, where even if the information is encrypted, by observing the sizes of packets, frequency of sending packets, and so on, an intruder can gain some information about the traffic that is being

exchanged. So, VPN defends against such traffic analysis attacks by aggregating the traffic from different sources and by aggregating that traffic into one flow that is sent between the gateway routers. So, hence, a VPN is secure. Apart from providing encryption, message integrity, and authentication, it also defends against traffic analysis attacks. So, in summary, we discussed IPsec and virtual private networks.

We discussed the ESP protocol for encapsulating packets that are sent over a security association, and then we discussed virtual private networks, which are used to provide the appearance of a private network but over the public internet. In the next lecture, we'll discuss how the internet key exchange protocol works. Recall that the IKE protocol is used to set up SAs. So, we'll discuss the operation of the IKE protocol in the next lecture. Thank you.