Network Security

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

Week - 06
Lecture - 32

IPsec and Virtual Private Networks (VPNs) for Network-Layer Security: Part 1

Hello, in this lecture and the next two lectures, we will discuss a protocol called IPsec, which is used to achieve network layer security. And we'll also discuss virtual private networks. As users, we are familiar with virtual private networks. When we are outside the institute or company campus, we often connect to the campus network using a virtual private network. So, virtual private networks are implemented using IPsec.

We'll discuss VPNs as well in these lectures. So, first we provide an overview of IPsec. So, IPsec can be used to establish a secure connection between the network layer entities at two nodes. These nodes may be hosts or routers. Recall the protocol stack.

We have the application, transport, network, link, and physical layer. Earlier, we discussed the protocol SSL or TLS, which provides security at the transport layer. Now, we discuss IPsec, which provides network layer security. IPsec provides the following security services. One is confidentiality.

So, confidentiality is achieved by encrypting the payload of IPsec. So, the payload of IPsec is the information that is provided by the transport layer and application layer. So, the payload of IPsec contains the transport layer header and application layer data. So, this payload is encrypted to achieve confidentiality. IPsec also provides endpoint authentication.

That is, suppose Alice and Bob are the network layer entities at the two sides of the IPsec connection. Then, Alice can prove to Bob that she is indeed Alice, and Bob can prove to Alice that he is indeed Bob. So, endpoint authentication is provided by IPsec. Message integrity is also provided, so we'll see that in every packet there is a message authentication code. So, this can be used to verify whether the source of the packet is legitimate and whether the packet has been tampered with during transit.

So, message integrity is also provided by IPsec. We discussed earlier some attacks in the case where an intruder intercepts the path between the transmitter and receiver. The intruder can replay packets, delete some packets, or reorder packets. So, IPsec provides mechanisms whereby the receiver can detect such attacks. The receiver can detect duplicate packets that the attacker may insert, and deletion or reordering of packets.

So, IPsec provides all these security services. We'll discuss how IPsec provides these services. We now provide an overview of virtual private networks (VPNs). So, virtual private networks are often created using IPsec. So, to understand VPNs, assume that a company has offices at multiple locations, say one office in Mumbai, one office in Bangalore, one in Hyderabad, and so on.

And the company wants to securely connect together all the machines in all the offices, via the public internet. So, in each of these campus networks, the campus in Mumbai, the campus in Hyderabad, and the campus in Bangalore, there are a large number of computers, maybe hundreds of computers in each office. And the company wants to connect together all the machines in all the offices via the public internet. So, in this figure, an example is shown. This is the public internet, and there is one company which has headquarters over here and a branch office.

For example, the headquarters might be in Mumbai, and there might be a branch office in Bangalore. There might be another branch office in Hyderabad, say. The company also wants to connect end systems of employees who are not currently in the office. For example, salespersons who are traveling away from their office and employees who are working from home to the office network. So, an example is shown here.

There is a salesperson in a hotel, and the salesperson wants to connect to the company network over the public internet. So, a virtual private network can be used to achieve this. So, the virtual private network provides connectivity in such an instance over the public internet. So, the transmission medium is the public internet, over which these computers in different offices and the salespersons exchange messages over the transmission medium, which is the public internet. But from the user's point of view, the network is just like a private network using leased lines.

From the point of view of the employees who are working in the different offices, such as Mumbai, Hyderabad, and Bangalore, and for the salesperson, it's just as if they are all on the same local area network. So, we'll see that packets are tunneled over the public internet. So, for example, there is a tunnel created between this gateway router of this branch office and this gateway router of this headquarters office. So, whenever a packet is sent from one computer over here to one computer over here, then that packet is tunneled over the public internet to this gateway router, and then the packet is extracted and just sent to this computer over here. So, we'll discuss how that tunneling is done using IPsec.

So, in summary, a virtual private network provides the appearance of a private network but uses the public internet as the transmission medium. We now provide an overview of the operation of IPsec. Suppose we want to establish a secure connection between the network layer entities at two nodes, say R1 and R2. So, in this case, two network layer logical connections, one in each direction, are established over which data can be securely exchanged. And these connections are known as security associations (SA).

So, it's important to note that there is one SA in each direction. So, there is one logical connection in each direction, one from R1 to R2 and one from R2 to R1. So, each SA is unidirectional. It transfers data only in one direction, either from R1 to R2 or from R2 to R1. So, for bidirectional data transfer, we require two SAs to be set up, one from R1 to R2 and one from R2 to R1.

To establish an SA from router R1 to router R2 or from R2 to R1, first, the two ends need to authenticate each other and also agree on the encryption and message integrity algorithms and keys. So, this is done by the routers R1 and R2 by executing a protocol known as the Internet Key Exchange Protocol (IKE). So, this IKE protocol is executed at the beginning to authenticate each other and to agree on the different algorithms and keys that will be used during the actual IPsec execution. Subsequently, after this Internet Key Exchange Protocol executes, the security associations are created, one in each direction, and subsequently, using IPsec, packets are transferred over these logical connections that have been set up. So, the transfer of packets over these security associations is secure.

Now, we discuss a security association (SA). IPsec packets are sent between the network layer entities at two nodes. They may be hosts or routers. Before sending IPsec packets from a source entity to a destination entity, a network layer logical connection called an SA is established. So, we discussed that during the execution of the Internet Key Exchange Protocol, SAs are established, one SA in each direction, and after that, IPsec packets can be sent from the source to the destination over this secure logical connection called SA.

And a security association is simplex, that is, it is unidirectional from the source to the destination. So, this allows flexibility. If we want to only send packets from R1 to R2, then we can set up an SA from R1 to R2. So, if we want bidirectional traffic, which is the more common case, in that case, we can always set up one SA in each direction. If both entities want to send secure packets to each other, then two SAs, one in each direction, need to be established.

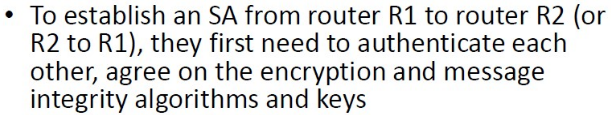

- To establish an SA from router R1 to router R2 (or R2 to R1), they first need to authenticate each other, agree on the encryption and message integrity algorithms and keys

So, here’s an example which illustrates security associations. Consider this corporate virtual private network, where there’s a branch office, there’s a headquarters office, and there are several salespersons who are working away from each of these offices. So, in particular, there are n traveling salespersons who are all at different locations, and these locations are different from the branch office and headquarters. Suppose there is bidirectional IPsec traffic between the headquarters and the branch office and between the headquarters and each of the salespersons. So, between this branch office and the headquarters, there is bidirectional traffic.

So, traffic is sent from this router to this router and from this router to this router over the public internet. Similarly, between the headquarters and each of the salespersons, there is bidirectional traffic. So, for example, there is bidirectional traffic from this router to this salesperson and from this salesperson to this router. So, what is the total number of security associations that are required in this case? So, my claim is that the total number is 2n+2.

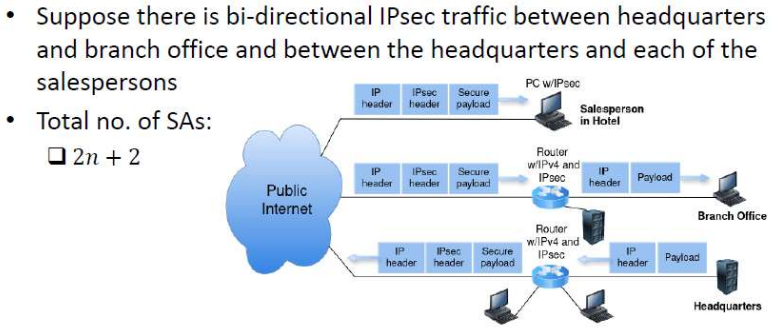

- Suppose there is bi-directional IPsec traffic between headquarters and branch office and between the headquarters and each of the salespersons
- Total no. of SAs:
  ❑ $2n + 2$

It’s easy to see why. For bidirectional traffic exchange between this router and this router, that is the headquarters gateway router and the branch office gateway router, we require two security associations, one in each direction. So that accounts for this term 2. These are the connections between the branch office gateway router and the headquarters gateway router. And we need two connections between the headquarters gateway router and each salesperson.

So, for example, between the headquarters gateway router and this salesperson, there will be two security associations, one in each direction. So, 2 security associations per salesperson, hence since there are n salespersons, a total of 2n security associations between the headquarters gateway router and the salespersons. Hence, the total number of security associations required is 2n+2. So, in such an example, where a virtual private network is set up, it's important to note that not all traffic sent into the internet by the gateway routers or by the salesperson's laptops will be secured using IPsec. For example, consider a computer somewhere in the headquarters office.

If this host accesses a web server that is owned by Amazon or Google, then this computer would use SSL, which we discussed earlier. They wouldn't use IPsec. So, IPsec is only used to create these tunnels between the headquarters gateway router and the branch office gateway router, as well as the tunnels between the headquarters gateway router and the salesperson's office. But if there is some traffic which is not flowing among these entities, salesperson's branch office and headquarters, then some other protocol like SSL will be typically used. For example, if there's a computer in the headquarters office communicating with some host somewhere in the public internet, over here, in that case the protocol SSL will be used for that communication.

We won't use IPsec. Consider this scenario where there is an SA from router R1 to router R2, as shown in this picture here. This is the headquarters office, and this is the branch office. And this is the gateway. R1 is the gateway router of the headquarters location, and R2 is the gateway router of the branch office. And this is the public internet.

And there is one SA from router R1 to router R2. So, routers R1 and R2 maintain the following state information about this security association. So, this state information tells R1 and R2 how to encrypt packets that are sent over this SA and how to provide message integrity in those messages. This information is provided to these entities, R1 and R2, by this state information. So, the following state information about this security association is maintained by R1 and R2.

One is a 32-bit identifier called the Security Parameter Index (SPI). So, at R1, the pair (SPI, destination IP address), that is the IP address of R2, of R2's interface, that is this one, this pair (SPI, destination IP address) uniquely identifies an SA. So, for example, there might be multiple security associations from R1 to R2. There may be other security associations from R1 to other routers. For example, there may be multiple branch offices.

So, there may be security associations from R1 to the gateway routers of the other branch offices as well. So, to uniquely identify each security association, we use this identifier, the security parameter index, along with the destination IP address. Now, the other state information about this SA that is stored at R1 and R2 is the origin interface of the SA and the destination interface. That is, the IP addresses of the origin interface and the destination interface. For example, the IP address of the origin interface is 200.168.1.100.

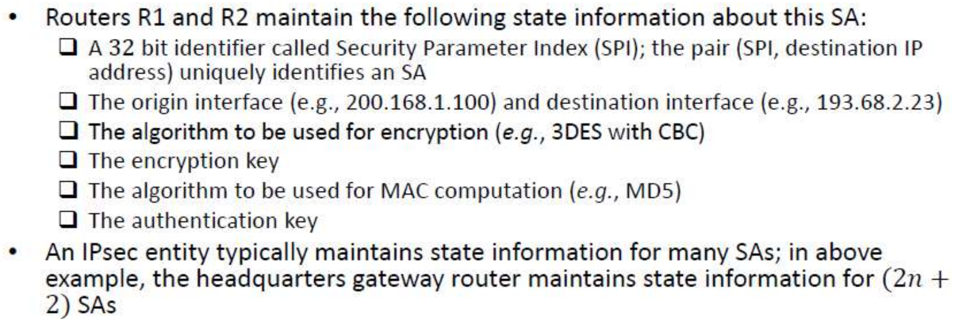

- Routers R1 and R2 maintain the following state information about this SA:
  - □ A 32 bit identifier called Security Parameter Index (SPI); the pair (SPI, destination IP address) uniquely identifies an SA
  - □ The origin interface (e.g., 200.168.1.100) and destination interface (e.g., 193.68.2.23)
  - □ The algorithm to be used for encryption (*e.g.*, 3DES with CBC)
  - □ The encryption key
  - □ The algorithm to be used for MAC computation (*e.g.*, MD5)
  - □ The authentication key
- An IPsec entity typically maintains state information for many SAs; in above example, the headquarters gateway router maintains state information for $(2n + 2)$ SAs

And the IP address of the destination interface is 193.68.2.23 in this example. So, the state information is which algorithm is to be used for encryption of packets that will be sent over this security association. An example algorithm might be 3DES with cipher block chaining. So, which algorithm will be used for encryption is part of the state information about this SA. Then, what is the encryption key?

That is, the secret key used for encryption of the packets sent from R1 to R2 over this SA. So, that is also part of the state information. Similarly, the algorithm to be used for MAC computation, for example, MD5 or SHA-1, and so on, which algorithm is to be used for MAC computation is part of the state information and the authentication key. Recall that for message m, the MAC is H(m,s), where s is a secret key known as the authentication key. So, this authentication key for this SA is also part of the state information.

So, in summary, routers R1 and R2 maintain all this state information about the SA from R1 to R2. An IPsec entity such as R1 or R2 typically maintains state information for many SAs. So, in this example, the headquarters gateway router, which is this one, maintains state information for 2n+2 security associations. So, we discussed in the previous example that there were 2n+2 security associations. In each of these SAs, this headquarters gateway router is either the source or the destination.

So, the headquarters gateway router maintains state information for 2n+2 security associations in that example. All this state information is stored in a database called the security association database (SAD). So, all this information for each of the security associations that are incident at a particular router is stored in this security association database at the router. So, whenever router R1 needs to construct an IPsec packet for forwarding over this SA, it accesses the state information of the SA in the security association database to find out how it should encrypt the packet and compute the message authentication code. So, when R1 needs to send a packet over this SA, it will look up the state information to find out which encryption algorithm it should use, what is the encryption key and what is the MAC algorithm and what is the authentication key.

So, all the state information will be looked up by R1 in the security association database. Similarly, router R2 has to receive packets over this security association. So, it uses its state information in its local security association database to find out how it should decrypt the packet and verify the message authentication code for a packet that is received over the security association. So, state information is stored at R1 and R2 and is looked up by R1 and R2 whenever they want to communicate using this security association. Now, to establish a security association from router R1 to router R2, they first need to authenticate each other and they need to agree on the encryption and message integrity algorithms and the keys as well as the security parameter index.

So, here the scenario is similar to what we had for SSL. So, there also, there was a handshake phase where the two sides authenticated each other and agreed on the encryption and message identity algorithms and the keys. So, we have a similar scenario here. So, before establishing the security association from R1 to R2, they need to authenticate each other and agree on the algorithms and keys to be used for encryption and message integrity. So, this is done using the IKE protocol.

Using the IKE protocol, the two sides can authenticate each other and agree on the encryption and message integrity algorithms, the keys, and the SPI. So, it has several similarities with the SSL handshake. So, we discussed that in the case of SSL, there are three phases. One is the handshake phase, then the key derivation phase, and the data transfer phase. So, in the SSL handshake phase, all these operations are performed: authentication and agreement on the algorithms and keys.

So, the IKE protocol has the following similarities with the SSL handshake. As in the SSL handshake, the two entities exchange certificates, which contain their public keys for authentication. And the two entities also negotiate the encryption and MAC computation algorithms. In particular, one entity, say R1, sends the list of supported encryption and

MAC computation algorithms to R2. And then R2 selects one encryption algorithm and one MAC computation algorithm from the list.

And then subsequently, these are used for the encryption and message integrity of the messages sent over the SA. So, this negotiation of encryption and MAC computation algorithms is done as part of the IKE protocol. And as part of the IKE protocol, R1 and R2 also exchange key material from which the encryption and authentication keys are derived. Recall that in the SSL handshake, the two sides exchanged the pre-master secret from which the keys were subsequently derived. So similarly, in the IKE protocol as well, some key material is exchanged, which is similar to the pre-master secret for SSL, from which the encryption and authentication keys will be derived later.

So, the IKE protocol has all these similarities with the SSL handshake. But there are some differences between the IKE protocol and the SSL handshake. We'll discuss the IKE protocol in detail later on, at which time these differences will become clear. So, in summary, we are discussing IPsec and virtual private networks. We provided an overview of IPsec and virtual private networks, and we discussed the concept of security association and also the functions of the IKE protocol.

We'll continue this discussion in the next lecture. Thank you.