Network Security
Professor Gaurav S. Kasbekar
Department of Electrical Engineering
Indian Institute of Technology, Bombay
Week - 07
Lecture - 42
Securing Wireless LANs : Part 8

Hello, recall that in the previous lecture, we discussed different types of management frames in 802.11. First, we'll discuss different types of attacks pertaining to management frames, and then we'll discuss how the 802.11w standard can be used to secure management frames. So, one type of attack on management frames is a deauthentication attack. So, recall that the deauthenticate frame allows clients and access points to request deauthentication from one another. Now, before 802.11w, no message integrity mechanism was used for this frame.

That is, suppose a client sends a deauthenticate frame to the access point. So, the access point has no way to find out whether this was indeed sent by the client or it was sent by someone else. Also, there was no way to detect any tampering in this deauthenticate message. So, consequently, the attacker could spoof this frame, either pretending to be the access point or the client, and direct it to the other party. So, this figure shows an example.

This is a client, and this is the access point, and this is an attacker. The vertical axis is the time axis. First, the authentication process runs. So, the client sends an authentication request to the access point, and then the access point responds with an authentication response frame. So, at this point, the client is authenticated with the access point.

Then, the client sends an association request to the access point, and the access point responds with an association response. At this point, the client is associated with the access point. Now, there is an attacker who wants to attack this network. So, the attacker sends a deauthentication frame to the access point, pretending to be the client. So, this deauthentication frame reverses these processes, and it deauthenticates the client from the access point.

So, once this deauthentication frame is received by the access point, the client gets deauthenticated from the access point. Next, the genuine client sends a data packet to the access point. The client is not aware that it has already been deauthenticated. So, the client sends its data packet to the access point, but then in the access point’s memory, the client is already deauthenticated. So, the access point responds with a deauthentication message.

So, the client is not able to communicate with the access point. This results in a denial of service to the client. So, this attack is possible because there is no message integrity in the deauthentication frame. The access point is not able to check whether this frame was sent by the genuine client or by an attacker. Because of this, the attacker can send a deauthentication frame and deauthenticate the client from the access point.

This is the deauthentication attack. When the deauthentication frame is sent by the attacker, the recipient of the frame, that is the access point of the client, would exit the authenticated state and refuse all further packets until authentication was re-established. So, that’s what happened here. When a data frame is sent by the client to the access point, it refuses that data packet and only responds to the deauthentication frame. So, it keeps refusing data packets until the client re-authenticates and reassociates itself.

So, if this attack is repeated persistently, the client may be kept from transmitting or receiving data indefinitely. So, the client is initially authenticated and associated, and then it sends some data to the access point. But then the attacker sends a deauthentication frame to the access point, and this deauthenticates the client. Suppose the client again performs these processes, authenticates itself, and associates with the access point. Then again, the attacker may send the deauthentication frame to the access point, and again the client will get deauthenticated.

So, this process keeps on repeating, and hence the client wastes much of its time just performing these authentication and association tasks. So, there is not much time available for actually exchanging data packets. So, hence, this is an example of a denial of service attack. By repeatedly sending deauthentication packets, the attacker can prevent the communication of legitimate data packets between the client and the access point. Another attack that is similar to the deauthentication attack is the disassociation attack.

It’s similar to the deauthentication attack, which is shown in this figure. This is the same figure as on the previous slide. So, this shows the deauthentication attack. And a disassociation attack is similar. Recall that before exchanging data packets, the client associates with an access point.

So, that is done through the exchange of association request and association response packets between the client and the access point. Now, 802.11 provides a disassociation message that is similar to the deauthentication message that we discussed earlier. So, this disassociation message ends the association between the client and the access point. Earlier, suppose the client and the access point were in this state where the client was associated with the access point. Then if a disassociation message is sent by the client to the access point or the access point to the client, it reverses the association and moves the client back to the authenticated state where it is authenticated but not associated.

So, it ends the association between the client and the access point. Before 802.11w, no message integrity mechanism was used for the disassociation message. So, hence, an attacker could send a disassociation message to either the client or the access point and the association between the client and the access point. This vulnerability can be exploited by an attacker by sending a disassociation message to the client or the access point. In response, the recipient of the message, that is, the client or the access point, would exit the associated state and would refuse all further packets until association was re-established.

This packet, which is the disassociation packet, is also similar to the deauthentication packet. So, it ends the association between the client and access point, and hence they are not able to exchange data packets immediately. The client has to associate with the access point again, and only after it associates again, it is able to exchange data packets with the access point. If we compare the dissociation attack and the deauthentication attack, we can see that the disassociation attack is less efficient from the point of view of the attacker than the deauthentication attack. If the attacker sends a deauthentication packet to one of the parties, client or access point, then they not only get disassociated but also get deauthenticated.

So, to exchange data packets again, the client and access point have to first authenticate via these packets, and then they have to associate. Only then can they start exchanging data packets. Whereas if an attacker sends a disassociation packet to either the client or the access point, then they get disassociated but they are still authenticated. So, to be able to exchange data packets, they only have to associate with each other. That is, the client has to associate with the access point, and then they can start exchanging data.

Hence, the deauthentication attack is more powerful than the disassociation attack because, to get back to the state where they can exchange data packets, the client and access point have to re-authenticate as well as re-associate. Only then can they start exchanging data.

So, to summarize, the deauthentication frame forces the victim to do more work to return to the associated state than does disassociation, ultimately requiring less work on the part of the attacker. So, the attacker has to send deauthentication frames less often than it has to send disassociation frames. So, the deauthentication attack is more effective from the point of view of an attacker than the disassociation attack.

There are also attacks on the power-saving functions that can be performed. So, first, let's understand power-saving functions in 802.11, and then we'll discuss attacks on these power-saving functions. So, power-saving functions also involve management frames. So, to conserve energy, clients are allowed to enter a sleep state during which they are unable to transmit or receive packets. So, in the sleep state, the wireless transceiver of a client is switched off.

So, hence it is not able to receive packets or transmit packets. So, this is done to conserve energy at the client because the client is typically a battery-operated device, such as a smartphone or a laptop. So, to conserve energy, the client goes into the sleep state when it is not going to exchange data packets in the near future. Before entering the sleep state, the client announces its intention so the access point can start buffering any inbound traffic for the client. So, the client may be in sleep state, but some node on the internet may send a packet to the client.

The access point is not able to immediately send the packet to the client because the client is in sleep state. So, the access point buffers the traffic that is being sent from nodes on the internet to the client. So, for the access point to know that it has to buffer traffic for the client, first the client has to inform the access point that it is going to transition to the sleep state. Occasionally, the client awakens and polls the access point for any pending traffic. So, when the client pulls the access point for any pending traffic, if there is any buffered data at the access point, which is sent by some node in the internet to the client, then the access point delivers the buffered data and subsequently discards the contents of its buffer.

So, when the client pulls the access point, the access point delivers the data that it had buffered when the client was in the sleep state. And once this data is delivered to the client, the access point deletes that data from its buffer. Now, before 802.11w, by spoofing the polling message on behalf of the client, an attacker could cause the access point to discard the client's packets while it was asleep. So, the client was still in the sleep state, but the attacker spoofed the polling message, pretending to be the client, and sent the polling message to the access point. This caused the access point to send the buffered packets to the client and then immediately discard the buffered packets after they were sent.

So, because of this, when the legitimate client wakes up from its sleep state, it does not get the previously buffered packets because they have been discarded by the access point. So, this was one attack that was possible again because this management frame, which is the polling message, was not protected. Another attack was that it was potentially possible to trick the client node into thinking that there were no buffered packets at the access point when in fact there were. So, in particular, the presence of buffered packets is indicated in the periodically broadcast packet called the Traffic Indication Map (TIM). The periodically broadcast packet is called the Traffic Indication Map (TIM).

So, this indicates the presence of buffered packets. So, there is a bitmap in the TIM that indicates whether there is buffered data for a particular client or not. So, this bitmap indicates the availability of buffered data for each client, which is associated with the access point. So, this bitmap is included in the Traffic Indication Map. Now, if the TIM message itself is spoofed, then an attacker may convince a client that there is no pending data for it, and the client will immediately revert back to the sleep state.

So, an attacker may spoof a TIM message and flip the bit, which indicates whether or not there is buffered data for a particular client. So, since this bit is flipped, the client may feel that there is no data buffered for it at the access point. So, because of the spoofing of this TIM message, the client will think that there is no pending data for it, and it will immediately revert to the sleep state, and the buffered packets will not be transferred from the access point to the client. Finally, the power conservation mechanisms rely on time synchronization between the access point and its clients so the clients know when to awake. So, the local clocks at the access point and the clients have to be synchronized with each other so that the clients wake up at the correct instance in order to receive the buffered packets.

But key synchronization information such as the period of TIM packets and a timestamp that is broadcast by the access point were sent without message integrity and in the clear. So, by forging these management packets, an attacker could cause a client node to fall out of synchronization with the access point. And hence the client would fail to wake up at the appropriate times. By compromising this synchronization between the client and the access point, the client is not able to wake up at the correct instance to receive the buffered packets from the access point. Because of this attack on the synchronization between the client and the access point, the client would fail to wake up at the appropriate times and receive the buffered packets from the access point.

So, this was another attack that was possible, again because management frames were not protected. So, now we discuss 802.11w and how it protects management frames. So, recall that before 802.11w only data frames could be protected in Wi-Fi and management frames were sent without any protection. And because of this, Wi-Fi had several vulnerabilities. Now the 802.11w amendment, which was introduced in 2009, provides protection for some management frames.

So, how does it provide protection for management frames? It uses existing security mechanisms, those which are standardized in 802.11i and which we discussed earlier. It uses these existing security mechanisms rather than creating new security schemes or new management frame formats. So, these existing security mechanisms are used to protect management frames in addition to data frames. Only some of the management frames are protected by 802.11w.

So, it is infeasible to protect the management frames that are sent before the four-way handshake because they are sent prior to key establishment. These management frames, which are sent before the four-way handshake, cannot be protected because at that point the client and access point do not have any secret key with which to protect the management frames. So, hence, it is not feasible to protect these management frames sent before the four-way handshake. The management frames that are sent after key establishment are protected. So, once the four-way handshake takes place, the client and the access point have secret keys, which they can then use to protect the subsequent management frames.

Some examples of management frames that are not protected are as follows: Beacon and probe request response frames. So, note that a beacon is used by clients to discover the presence of access points in the vicinity. So, at that point, the client has not associated with the access point. So, they have also not done the four-way handshake.

So, clearly, the beacon frame cannot be protected using 802.11w. Similarly, the probe request and response frames are sent before they are authenticated and associated. So, hence, the probe request and response frames can also not be protected. Then, authentication request and response frames, and association request and response frames. These are also exchanged before the four-way handshake.

So, at that point, the client and access point don't have any secret keys. Hence, these frames are also not protected. Examples of management frames that are protected by 802.11w are disassociate frames and deauthenticate frames. So, after the association request and response frames are exchanged, then the four-way handshake takes place, and after this, if any disassociate frame or deauthenticate frame is sent, then it is protected using 802.11w. Now, protection-capable management frames are protected using the same mechanism as an ordinary data frame, which we discussed earlier.

In particular, the payload is encrypted using CCMP, which we discussed earlier, and message integrity is provided for the payload and the header using CCMP as well. So, we discussed how encryption and message integrity can be provided using CCMP. So, we discussed that for data frames, and exactly the same mechanisms are used for protecting management frames as well. So, the same procedure is used as in 802.11i, namely CCMP, but that is used for protecting management frames instead of protecting data frames. As for data frames, the temporal key or TK is used to provide encryption and message integrity for management frames.

Replay protection is provided by using the same mechanism as for data frames. So, if an intruder takes an old frame and replaces it, then that can be detected using the mechanism for replay protection that we discussed earlier. So, the mechanism used for replay protection in data frames is also used for replay protection of management frames. Now, in 2018, WPA3, which stands for Wi-Fi Protected Access 3, was announced as a replacement for WPA2. Recall that WPA2 is the same as 802.11i.

WPA2 was announced in 2004, but then in 2018, WPA3 was announced as a replacement for WPA2. WPA3 makes it mandatory to use management frame protection that is defined in the 802.11w standard. So, WPA3-compatible devices use 802.11w to protect the management frames that are exchanged in the network. As an aside, WPA3 also makes it mandatory to use a protocol called Simultaneous Authentication of Equals (SAE), also known as the dragonfly handshake. This is used for authentication between the client and the access point.

So, WPA3 introduces this handshake called the dragonfly handshake, or SAE, for authentication between the client and access point. We will not get into the details, but you can read about this exchange, SAE, or dragonfly handshake if you're interested. So, due to management frame protection, the attacks that we discussed above are defended against. So, for example, suppose a client and access point are already associated with each other and they have also done the four-way handshake. Subsequently, if an attacker tries to send a deauthenticate packet to the access point or the client, pretending to be the other party, in that case, that will be detected because that management frame won't be protected.

So, because of the protection of these management frames, attacks such as deauthentication attack, disassociation attack, these can be defended against. But there are vulnerabilities even in 802.11w networks. So, recall that 802.11w provides protection to several management frames. But even in 802.11w networks, some vulnerabilities have been shown to exist. We now discuss some of them.

So, one is the deauthentication attack. 802.11w provides protection to deauthentication frames only after the completion of the four-way handshake. So, that's because before the four-way handshake, the client and access point don't have any secret keys using which the deauthentication frames can be exchanged, using which the deauthentication frames can be protected. So, because of this, 802.11w only provides protection to deauthentication frames after the completion of the four-way handshake. So, because of this, an attacker can send a spoofed deauthentication packet before the four-way handshake has completed.

So, this will end the connection between the client and the access point. So, suppose the client and access point have authenticated, and the client has associated with the access point. At this point, the attacker can send a spoofed deauthentication packet, and that will end the connection between the client and access point. The client will have to authenticate again and then associate again with the access point and then complete the four-way handshake. So, this is an example of a deauthentication attack.

If this attack is performed repeatedly, it results in a denial of service to the client. So, the client authenticates itself and then associates with the access point, but before it is able to complete the four-way handshake, this attack takes place—the deauthentication attack—and hence it goes back to the deauthenticated state, and then it has to re-authenticate and re-associate. So, if this is performed repeatedly, it results in a denial of service to the client. Another example of an attack is a beacon probe frame flood attack. So, recall that 802.11w does not protect beacons.

And that's because beacons are sent right at the beginning before the four-way handshake is completed. So, an attacker can flood the network with a large number of beacons which advertise different access points. So, legitimate beacons are sent by legitimate access points, but the attacker can flood the network with a large number of fake beacon packets which advertise different access points which are not actually there in the vicinity. So, this kind of flooding will confuse the clients. This confuses clients and makes it difficult to find the legitimate access point.

Because of the beacon flood attack, apart from the genuine beacon packets that advertise the legitimate access points, there will also be several fake beacon packets, and it will be difficult for clients to distinguish the legitimate beacon packets from the fake ones. Similarly, recall that 802.11w does not protect probe requests and responses. Again, an attacker can flood the network with a large number of probe requests or response frames, and this will again confuse clients and make it difficult to find the legitimate access point. So, these attacks on beacon frames and probe request response frames are possible because they are not protected using management frame protection. So, these kinds of attacks can waste bandwidth and/or cause confusion to clients and/or access points.

So, attacks like these can be defended against using systems known as intrusion detection systems. Later on, we'll discuss firewalls and intrusion detection systems. So, these intrusion detection systems can be used to defend against these kinds of attacks. So, in summary, we discussed different management frames in 802.11, and then we discussed different types of attacks that are possible on these networks. For example, deauthentication attacks and disassociation attacks.

And then we discussed 802.11w, which provides protection to management frames and can be used to defend against several of these attacks. But we also noted that 802.11w does not defend against all attacks because management frames that are exchanged before the four-way handshake are not protected using 802.11w. So, hence, there are some attacks still possible, which can be defended against using intrusion detection systems. So, this concludes our discussion of the security of wireless local area networks, in particular, Wi-Fi security. Thank you.