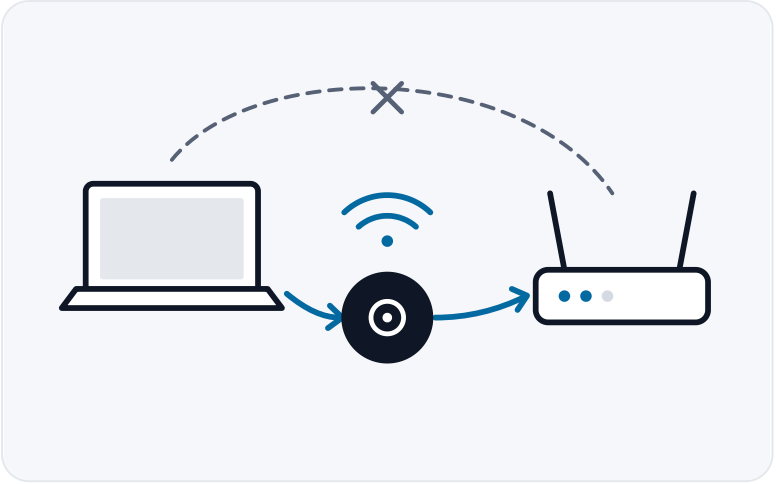
A-SIT Technology

E-Government Innovation Center

12 January 2026 Analysis · Web & Cloud · Mobile Security · 23 min read

Man-In-The-Middle Attacks against Wi-Fi

Florian Draschbacher



Authors	Florian Draschbacher
Version	2026-01-07
Source	https://technology2.a-sit.at/en/publications/man-in-the-middle-attacks-against-wi-fi/
Note	Machine translation.

Key findings

- ARP spoofing, DHCP spoofing and multi-channel MITM still work against Samsung and Google Android devices running Android 14 to 16.
- In DHCP spoofing, the Android DHCP client ignores NAK replies, so the attacker can register itself as the default gateway.
- The multi-channel attack did not work reliably in the test setup; in some cases clients did not switch frequency.
- WPA3 does not protect fully either: beacon frames are exempt from Management Frame Protection, so Channel Switch Announcements can still be forged.
- On public networks, end users should check that application-level encryption (HTTPS) is used and should rely on a trustworthy VPN.

Wireless network communication via Wi-Fi has today become a taken-for-granted part of everyday life. Especially on mobile devices such as smartphones, the technology is used in a wide variety of ways—for example, not only to connect the device itself to the Internet, but also to share mobile network connections with other devices. Despite the commonplace use of this technology, most users are not aware of the many security weaknesses that have in some cases been known to experts for decades.

Within the scope of this project, the concrete scenario of Internet sharing via Wi-Fi tethering was used to investigate to what extent the network stacks of current smartphones are affected by known Wi-Fi security weaknesses, and which attacks are enabled as a result. In this report, we describe three different attacks and their implementation, evaluate them in the given scenario, and compare their characteristics. Finally, we discuss possibilities for detecting and preventing these attacks.

1. Introduction

The wide availability of fast wireless Internet via Wi-Fi is one of the cornerstones of today's digital world. Modern mobile phone contracts now offer such large data volumes that smartphones can be online permanently, but end devices such as laptops still only exceptionally have mobile network modules. If a laptop is to be connected to the Internet while on the move, Wi-Fi is usually used. Hotels, restaurants, cafés or public transport increasingly offer free Wi-Fi hotspots for this purpose. If none is available, users usually resort to Wi-Fi tethering. In this case, the smartphone's mobile-network-based Internet connection is passed on to the laptop via Wi-Fi. On a technical level, the smartphone starts a Wi-Fi hotspot for this purpose and bridges data between the Wi-Fi and the mobile network adapter.

The wireless transmission of network data via Wi-Fi enables usage scenarios that would not be possible with traditional cable connections, or only far less conveniently. For example, Wi-Fi allows the device to be moved freely within the reception range. While this gain in convenience is immediately apparent, wireless network connections also have different requirements with regard to transmission security.

The Internet and a large number of other network applications use the protocols of the TCP/IP family as the technical basis for communication. TCP/IP is also used for Wi-Fi connections to package the transmitted data and transfer it in an orderly manner between different devices. However, TCP/IP was originally designed for wired applications. In this configuration, threat scenarios are relatively manageable. Information is only exchanged along the physical connections, so a potential attacker first needs physical access to a network socket or a cable. In contrast, in Wi-Fi networks the information is (somewhat simplified) transmitted as radio waves in all directions. To prevent any bystander within radio range from reading or influencing the communication, the data must be encrypted. Another decisive difference results from the routing mechanisms used. In wired connections, modern network switches are usually used to connect several network participants with one another. These switches contain logic to forward network packets only to the recipient noted in the packet. An uninvolved device connected to the network can therefore not read communication exchanged between two other participants. In a Wi-Fi network, all participants are connected via a shared medium (the air, or radio waves in the air), so such filtering is not possible. A malicious network participant can therefore eavesdrop on the communication of all other

participants if no additional encryption is used. Regardless of the encryption, manipulation of the communication is possible. Specifically, in the simplest case, an attacker can, for example, send interference signals that cause the Wi-Fi packets not to arrive.

In view of the large number of possibilities for attacks on Wi-Fi compared to wired network connections, and the widespread use of Wi-Fi, especially in mobile devices, the question arises of the feasibility and practical relevance of attacks on Wi-Fi. In this project, the mobile operating system Android was therefore tested for its vulnerability to 3 different attacks. A particular focus was placed on attacks that allow the attacker not only to read but also to manipulate the decrypted communication. Although no fundamentally new attacks were found for this, existing approaches were adapted and optimised for the Android network stack. The goal of this report is to raise awareness of data security in Wi-Fi networks on the basis of the practical relevance of the various attacks. To this end, a brief overview of measures for detecting and preventing attacks is provided at the end. Both techniques that can be integrated into the existing protocol stacks and general recommendations for end users are discussed.

2. Background

This section explains the technologies that are necessary for understanding the later explanations.

2.1. TCP/IP

The Internet and numerous local network applications are based on the TCP/IP protocol family as the technical foundation for communication. TCP/IP defines a layer model with clear responsibilities: the link layer (physical medium and MAC addressing), the internet layer (IP addressing and routing), the transport layer (e.g. TCP, UDP) and the application layer (HTTP, DNS, among others). This division makes it possible for different transmission media such as copper cable, fibre optics or radio to use the same protocols above them.

At the link layer, Media Access Control (MAC) addresses uniquely identify network interfaces within a local transmission range (usually a Local Area Network, or LAN). MAC addresses are hardware-related and are used by network devices when sending network packets (so-called network frames) to specify destination and source within a LAN. At the internet layer, IP addresses (IPv4 or IPv6) take over the logical identification of hosts and are necessary so that routers can forward packets between different networks. Addressing in both layers is complementary: MAC controls local delivery, IP controls routing across networks.

The transport layer above distinguishes between connection-oriented services (Transport Control Protocol, or TCP) and connectionless services (User Data Protocol, or UDP). TCP ensures reliability through connection setup, sequencing, retransmission of lost packets and flow control - properties that are important for many applications. UDP remains lightweight and is suitable for time-critical applications, but is not reliable without additional measures.

2.2. ARP

As described above, network participants in TCP/IP networks usually have 2 different addresses: a MAC address and an IP address. The MAC address serves (at least in theory) to uniquely identify a network interface. It is usually assigned largely at random by the manufacturer at the factory. In the higher transport layer, the IP address is used to address packets. To be able to deliver packets of the transport layer

(addressing with an IP address) by means of frames on the link layer (MAC addresses), network participants must know under which MAC address the participant with the desired IP address is available. The Address Resolution Protocol (ARP) [1] is used for this. If a network participant wants to send a data packet to another participant with a certain IP address, it sends an ARP request by broadcast to all other participants in the network. The request contains the IP address whose owner is to be found, as well as the MAC address and IP address of the requesting party. Every participant that receives the request compares the contained IP address with its own. If there is a match, the participant replies directly to the requester. The latter can now begin delivering the IP communication.

Network applications are usually optimised with regard to performance characteristics such as data throughput and latency. Here, it is of course a hindrance if the MAC address must first be queried for every new connection at the IP level. For this reason, network participants store this mapping in a cache called the ARP table. ARP queries thus only need to be handled if there is no entry in the table for the desired participant (target host). To additionally ensure that ARP requests have to be sent as rarely as possible, participants usually proactively send so-called gratuitous (roughly: unsolicited) ARP messages when their IP address changes. These are ARP replies that are sent to all network participants in the LAN without a corresponding request. They are then taken into account in the ARP table.

2.3. DHCP

As a further important component of the TCP/IP network stack, DHCP, the Dynamic Host Configuration Protocol [2], must be described here. It is needed to dynamically provide the participants in a network with IP addresses. Once a network participant has connected to a network at link level, it sends a DHCP Discover message as a link-layer broadcast to all other participants. If a DHCP server is among them, it usually replies with a DHCP Offer, in which it offers the use of a specific IP address for a certain duration. If the new participant (DHCP client) wants to accept this offer, it replies with a DHCP Request, which is finally confirmed by the DHCP server with a DHCP Acknowledgement. To shorten the duration of this interaction, an accelerated process called Rapid Commit is available if server and client support it. Here, the DHCP server responds with a DHCP Acknowledgement already in its first reply instead of the DHCP Offer. In a single message, it thus already provides an immediately usable IP lease that the client no longer has to request separately.

DHCP is usually used not only for central and dynamic management of IP addresses, but also to tell new network participants the IP addresses of important points of contact in the network. The replies sent out by the DHCP server thus contain the IP addresses of the DNS server, which can be contacted to translate domain names into IP addresses, and of the default gateway. The latter serves as a bridge to other networks, and usually also to the Internet. Put simply, all packets that are to be sent to IP addresses that do not exist in the local network are handed to the default gateway for forwarding. It usually has routing tables that it can use to decide where to forward the packet. This mechanism is essential for the ability to transfer IP packets between different networks at link level.

2.4. Wi-Fi and WPA/WPA2/WPA3

Wi-Fi is based on the IEEE 802.11 standard and enables wireless network connections via radio waves. The standard defines various versions and extensions that offer different speeds, ranges and frequency bands. The most common frequency ranges are the 2.4 GHz and the 5 GHz band. The 802.11 standards comprise several variants, including 802.11a, 802.11b, 802.11g, 802.11n and 802.11ac, which enable increasingly higher

data transfer rates and improved network efficiency. WLAN is used worldwide in private households, offices and public areas to network devices such as computers, smartphones, tablets and smart home technology. Communication usually takes place between an access point and a client. In everyday use, ranges of up to 45 metres (2.4 GHz band) or 15 metres (5 GHz band) are possible. Depending on the environment and the end devices used, however, these values can be exceeded significantly.

To prevent communication from being eavesdropped on by attackers within this range, Wi-Fi networks are usually secured with a security protocol. In practice, WPA2 and WPA3 are usually used today. WPA2 was developed as the successor to WPA and WEP (Wired Equivalent Privacy), which are now classified as insecure. It is based on the AES (Advanced Encryption Standard) encryption algorithm. WPA2 protects both the authentication and the data transmission and ensures that only authorised devices can access the network. In many cases, networks use the WPA2-PSK (Pre-Shared Key) variant. This is based on a shared, predefined key with a minimum length of 8 characters. If a client wants to establish a connection with the access point, it communicates this request to the access point, which replies with a random value (Anonce). The client now also generates a random value (Snonce), which it sends to the access point. On the basis of the PSK, the network name (SSID; which the access point broadcasts unencrypted at regular intervals) and the two nonces, both communication partners then compute a Pairwise Transient Key (PTK), which is used to encrypt the further communication. Although WPA2-PSK enjoys great popularity, especially among private users, the standard contains decisive weaknesses. Not only do all network participants share the same pre-shared key (colloquially usually called the Wi-Fi password), but with WPA2-PSK the PTK can additionally be derived if the PSK and the client handshake are known. The client handshake can be trivially recorded by an attacker within Wi-Fi range. In many scenarios (such as public hotspots in cafés, hotels, etc.), the PSK is also known to a large number of people. The scenario of recording all transmitted data (that is not additionally encrypted at higher network layers) is therefore quite realistic. In contrast to WPA2, with the newer WPA3 standard the Pairwise Transient Key cannot be reconstructed even if the attacker knows the PSK and has a recording of the client handshake. WPA3 is therefore to be preferred over WPA2 in all use cases. Although many Wi-Fi hotspots already offer the possibility of WPA3 connections, they are often operated in transition mode to also ensure compatibility with older devices. The client then has the choice of whether WPA2 or WPA3 is used. If it chooses WPA2, the restrictions described above with regard to connection security still apply.

3. Concrete evaluation setup

In this project, the feasibility and effectiveness of various man-in-the-middle attacks on Wi-Fi connections of mobile end devices were evaluated. In a man-in-the-middle attack (or machine-in-the-middle attack), the attacker is in a position between two devices (referred to below as device A and B) that communicate with each other. Although it appears to the two devices as if they were communicating directly with each other, in reality all sent packets first go to the attacker, who then forwards them to the respective other device. Even if the packets are encrypted and the attacker has no knowledge of the key, the attacker thus has the possibility of influencing the communication, for example by selectively dropping packets, that is, not forwarding them to the other device. If the attacker knows the key (for example because WPA2 is used and the attacker knows the PSK and can thus derive the PTK from a recorded handshake) or if data is transmitted unencrypted (for example because no encryption at all is used in a Wi-Fi network), the attacker can read the data or even modify it before forwarding.

To be able to carry out such an attack in a real environment, the attacker must be within Wi-Fi range. For networks in the 2.4GHz spectrum, this means an approximate maximum distance from the hotspot of 30 metres, or 20 metres for networks in the 5Ghz spectrum.

In addition, the attacker must have suitable hardware, specifically a Wi-Fi interface that can be operated in monitor mode. Inexpensive variants can be purchased for around €20.

The following concrete scenario was evaluated for this project. A smartphone was used to start a Wi-Fi hotspot by means of tethering (that is, the possibility of wirelessly sharing the mobile Internet connection with other devices via Wi-Fi). This device also acts as device A, even though in general the functions of hotspot and device A do not necessarily have to be combined in the same device. A second smartphone was then used as a Wi-Fi client that connects to the Wi-Fi hotspot. This is the victim of the attack, referred to below as device B.

The attacker's goal in all cases was to redirect the network communication between devices A and B so that the attacker obtains a man-in-the-middle position. To this end, attack techniques at different levels of the network stack were evaluated. It must be noted that modern Android devices usually generate their MAC address randomly for each new network connection. The attacker therefore knows neither the MAC address nor the IP address of device B in advance.

All attacks were implemented on a Raspberry Pi 5 running Kali Linux. The attacks were each implemented in Python based on the Scapy library [3]. Since some attacks required a Wi-Fi interface that can be put into monitor mode, an external interface based on the Mediatek MT7921 chipset was connected via USB. In monitor mode, it can also read Wi-Fi frames that were addressed to other network participants. The effectiveness of some attacks can be increased by using several network interfaces. Here, we used 2 of the external interfaces described.

4. ARP spoofing

In ARP spoofing, the attacker tries to modify the ARP cache of the victim (device B) so that the resolution of a certain IP address (e.g. that of device A) results in the attacker's MAC address instead of the actual target device. To do this, the attacker sends gratuitous ARP messages to make other network participants believe that the attacker is the legitimate owner of the IP address of another network participant. Since the attack takes place above the link layer, the attacker must either find an open Wi-Fi network or know the PSK.

When implementing this attack, the attacker must first establish a connection to the Wi-Fi hotspot. It was further assumed that the attacker establishes this connection before the victim. The attacker must then detect as early as possible that device B is connecting to the network. The earlier this is detected, the shorter the time in which the ARP cache of device B contains the correct MAC address for device A. During this time, the attacker does not have a MITM position. For our implementation, we therefore observe the DHCP broadcasts in the network. As soon as device B has established the connection at link level, it requests an IP address by means of DHCP Discover. From this point on, the attacker has all the information necessary to start the actual ARP spoofing. The attacker now sends, in the quickest possible succession, gratuitous ARP messages that contain the attacker's actual MAC address as the source MAC address, but the IP address of device A. Once the victim has received the message and adapted its ARP table accordingly, it sends IP packets that were actually meant for device A to the attacker. The attacker can now decide whether to forward the packets to device A and whether to modify them beforehand.

To determine whether the attack works, an HTTP server was started on device A, which was accessed from device B. The attacker was implemented as a simple proxy that logs the received requests but forwards them unchanged to the actual target device.

In our evaluation, this attack worked reliably on various Android devices from Samsung and Google running Android 14, 15 and 16. It must be noted, however, that for a brief moment directly after device B establishes the Wi-Fi connection, its ARP table contains the correct MAC address of device A for its IP address. If device B establishes a connection to device A at exactly this moment, that connection is not affected by the MITM attack.

5. DHCP spoofing

DHCP spoofing also starts above the link layer, but somewhat earlier than ARP spoofing. The attack essentially consists of two steps. In the first step, the attacker claims all available DHCP leases for itself. In the second step, the attacker takes over the role of the DHCP server in the network and itself responds to DHCP requests. The actual goal of the attack is now not control over the assignment of IP addresses, but over the information on the DNS server and default gateway, which is also contained in the DHCP replies. If the attacker positions itself as the default gateway, then, for example, all packets that have to be forwarded to external networks (e.g. the Internet) go to the attacker.

To use up the server's DHCP leases, the attacker must send as many DHCP requests to the server as there are leases available. The exact number depends on the size of the subnet. The Android tethering implementation uses a subnet of at most 255 clients, so the number of required DHCP requests remains manageable. In the implementation, however, a few particularities must be taken into account, which will not be discussed in detail here so as not to provide assistance for real attacks. Our implementation is able to use up all leases in the tested scenario in less than a second.

For all further requests for an IP address, the DHCP server now replies with a NAK (Non-Acknowledgement), that is, it communicates that it unfortunately cannot be of service. The DHCP client implementation of Android (which mimics the behaviour of the Linux program dhcpcd) essentially ignores these NAK replies and keeps waiting for a positive DHCP reply including a DHCP lease. The attacker can take advantage of this peculiarity of the implementation: Since the client has no information about the DHCP server at the start of the interaction, it sends the DHCP Discover message as a broadcast to all participants in the network. The attacker therefore also receives it and can compose a reply based on the contained information (in particular the transaction identifier). In it, the attacker can enter arbitrary hosts as DNS server and default gateway. If the attacker designates itself, it now has a MITM position for all connections to external networks.

To determine whether the attack works, we assumed a somewhat modified scenario here for simplicity: We assumed that an application on device B queries the default gateway via the Android API and establishes an HTTP connection to it. We further assumed that the default gateway coincides with device A, as in the scenario described above.

This attack worked flawlessly on the mentioned Android devices from Samsung and Google running Android 14, 15, and 16.

6. Multi-channel man-in-the-middle attack

In contrast to the previously described attacks, this attack operates on the link layer, that is, on the Wi-Fi frames. The general mode of operation was first presented by Mathy Vanhoef et al. [4]. It is based on a gap in the Wi-Fi specification: As an unauthenticated attacker, it is possible to send network participants a request to change the communication frequency. The attacker does not need to know the PSK of the network for this. The hotspot itself does not notice this request. If a client carries out the switch, the attacker now acts as a proxy between the legitimate frequency of the hotspot and the different frequency to which the attacker has redirected the Wi-Fi clients. The attacker has thus effectively taken up a MITM position.

Like the other attacks, this attack was evaluated on smartphones from Samsung and Google with Android versions 14 to 16. Even though the attack works in principle, it must be mentioned that problems repeatedly occurred in the setup we tested. In some cases, the clients could not be persuaded to change the communication frequency and therefore continued to communicate directly with the hotspot. Vanhoef et al. recommend a slightly different setup for the attack, with Wi-Fi interfaces based on the Atheros AR9001 or AR9002 chipsets [5]. The problems with the reliability of the attack could be related to this difference.

7. Comparison of the attacks

All implemented attacks differ somewhat in terms of reliability and possible purposes. A brief comparison is therefore provided here.

The criteria used here for differentiation are:

- Functional level (in the network stack) of the attack This determines whether the attacker must establish a network connection before the attack. To do so, the attacker must either know the PSK or find an unencrypted network.
- Scope Which connections can be attacked?
- Reliability How reliably does the attack work, and what sources of unreliability are there?

The comparison between the evaluated attacks is shown in Table 1.

Attack	Functional level	Scope	Reliability
ARP spoofing	Internet layer (ARP)	Local connections	Very high after brief setup
DHCP spoofing	Internet layer (DHCP)	Connections to external networks	Very high
Multi-channel MITM	Link layer (802.11)	All connections	Depends on setup

Table 1 Comparison of the evaluated attacks

8. Countermeasures

As a result of our evaluation, we must conclude that all tested attacks can still be exploited effectively. In the following, we present, on the one hand, options for automatically detecting and preventing such attacks and, on the other hand, preventive measures that can protect specific users from an attack.

8.1. Automatic detection and prevention of the attacks

The fact that the tested attacks were never fixed despite being well known is due to the fact that it is often not easy to distinguish between malicious and benign behaviour. In the following, we nevertheless discuss options for how attacks can be detected or prevented at the protocol level.

- ARP spoofing In research, various heuristics have been presented in the past that try to detect ARP spoofing attacks either by passive observation of network communication [6] or by active probing [7]. If a network participant has been identified as malicious, it can be removed from the ARP table and any further communication can be ignored.
- DHCP spoofing The simplest solution would be for clients to no longer accept replies from other DHCP servers after receiving a NAK. Especially in the private user sector, the legitimate DHCP server is usually located on the host that acts as the Wi-Fi hotspot. Since in Wi-Fi networks communication between network participants must always go through the Wi-Fi hotspot, the legitimate DHCP server receives requests first and can also reply first. If every request contains a unique random transaction ID, there is no way for another network participant to reply before the legitimate DHCP server.
- Multi-channel MITM This attack only works as long as management frames for the channel switch are accepted unencrypted. A naive solution would therefore be to require encryption here. In principle, this is also implemented in WPA3, where Management Frame Protection must be used mandatorily. However, beacon frames are exempt from this rule, so Channel Switch Announcements can still be forged [8].

8.2. Preventive measures for end users

End users should exercise caution whenever they are in networks to which a large number of other (unknown) users have access. Especially in such networks, it is important to ensure that connections are encrypted at the application level. If the network is used for web browsing, the browser usually indicates if a page does not use TLS encryption. In the standard case (when using HTTP over TLS, i.e. HTTPS), there is then no danger that transmitted data can be spied on or modified. Nevertheless, an attacker can at least roughly trace from the contacted IP addresses which websites are being communicated with. This risk can be reduced further if users access the Internet only via a trustworthy VPN.

9. Summary

Wireless network connections via Wi-Fi are an everyday companion for most users. Despite this central role of the technology in the modern networked world, vulnerabilities in the protocols involved that can be exploited in attacks have been known for years. In this project, we investigated to what extent these attacks can be used against modern mobile devices in the concrete scenario of Wi-Fi tethering. We had to

conclude that mobile devices are still affected by ARP spoofing, DHCP spoofing and multi-channel man-in-the-middle attacks. In addition to a comparison of these attacks based on our implementations, in this report we also discussed options for preventing these attacks and for taking precautions against an attack.

References

[1] Network Working Group, „RFC 826," 11 1982. [Online]. Available: https://datatracker.ietf.org/doc/html/rfc826.

[2] Network Working Group, „RFC 1531," 1993. [Online]. Available: https://datatracker.ietf.org/doc/html/rfc1531.

[3] The Scapy community, „Scapy: the Python-based interactive packet manipulation program & library," 2024. [Online]. Available: https://github.com/secdev/scapy.

[4] M. a. P. F. Vanhoef, „Advanced Wi-Fi attacks using commodity hardware," in Proceedings of the 30th Annual Computer Security Applications Conference (ACSAC '14), 2014.

[5] M. Vanhoef, „Multi-Channel Machine-in-the-Middle," 2022. [Online]. Available: https://github.com/vanhoefm/mc-mitm?tab=readme-ov-file#3-launching-the-attack.

[6] A. I. A. E. B. Mohammed M. Alani, „ARP-PROBE: An ARP spoofing detector for Internet of Things networks using explainable deep learning," Internet of Things, Bd. 23, 2023.

[7] V. N. S. Ramachandran, „Detecting ARP Spoofing: An Active Technique," in Information Systems Security. ICISS 2005, 2005.

[8] M. Vanhoef, „Attacking WPA3: New Vulnerabilities & Exploit Framework," 2022. [Online]. Available: https://conference.hitb.org/hitbsecconf2022sin/materials/D1T1%20-%20Attacking%20WPA3%20%20New%20Vulnerabilities%20and%20Exploit%20Framework%20-%20Mathy%20Vanhoef.pdf.