Version 3 of the impact and preconditions of FragAttacks in *"Fragment and Forge: Breaking Wi-Fi Through Frame Aggregation and Fragmentation"*. When listing CVEs the prefix "CVE-2020-" is removed. A CVE may have several impacts and preconditions.

Vulnerability (= flaw)		Attack Objective		Preconditions					
CVE	Short Description	Impact	Target	attacker.com	Fragmentation	Rekeys	Predict IP ID	Other Vuln.	Additional Notes
24588	A-MSDU design flaw	Inject	Client / AP	●					
24587	Mixed key attack	Exfiltrate	AP	●	●	●			
24586	Fragment cache attack	Exfiltrate	Hotspot AP		●				
		Inject	Hotspot AP		●				Packets are injected under victim ID.
		Inject	Cautious client		●				Can inject packets in trusted network.
26145	Plain. broadcast fragment	Inject	Client / AP						
26144	A-MSDU EAPOL	Inject	Client / AP						
26140	Plaintext frames	Inject	Client / AP						
26143	Frag. plaintext frame	Inject	Client / AP						
26139	Forwarding EAPOL	Inject	AP					●	A client must be affected by 24588.
26146	Non-consec. PNs	Exfiltrate	AP	●	●				
26147	Only first frag. encrypted	Inject	Client / AP		●		●		Target must also be affected by 24588.
		Inject	Hotspot AP					●	AP must also be affected by 24586.
		Inject	Cautious client					●	Client must also be affected by 24586.
	Only last frag. encrypted	Inject	Client / AP		●				
26142	Fragment as full frame	Inject	Client / AP	●	●				
26141	No TKIP MIC check	Inject	Client / AP						Fragm. attacks possible with TKIP.

Impact This column shows whether a vulnerability can be exploited to inject or exfiltrate network packets (or both). In exfiltration attacks, a vulnerable AP (the target) is exploited to exfiltrate data sent by clients. If a client is vulnerable to packet injection, it can be tricked into using a malicious DNS server. If an AP is vulnerable to packet injection, the NAT/firewall might be bypassed (using "hole punching") to access devices in the local network. See the next page of details.

Target This column indicates whether, under the given preconditions in the row, clients or APs can be exploited (or both). Two special targets are:

- *AP Hotspot*: the flaw can only be exploited against APs in hotspot-type networks where users distrust each other.
- *Cautious client*: the flaw can only be exploited against clients that will connect to a network of which the adversary also knows the password (the client only sends sensitive data over *other* networks though). See Section 5.1 in the paper for details.

Preconditions This column lists the conditions that must hold for the flaw to be exploitable. Common conditions are:

- *attacker.com*: the client must connect to a server of the adversary. This can be as simple as tricking the client into downloading an image from the adversary's server. Note that (JavaScript) code execution is not required.
- *Fragmentation*: the peer must send fragmented frames.
- *Rekeys*: the network must be configured to periodically refresh the pairwise session keys of connected clients.
- *Predict IP ID*: the adversary must be able to predict the IP identification field values used by another peer.
- *Other Vuln.*: the target must be affected by a second vulnerability. See the column "additional notes" for details.

Summary The A-MSDU design flaw (highlighted in bold) may be exploited in practice in targeted attacks. The mixed key attack has numerous non-trivial preconditions and is therefore a theoretic attack. The cache attack is only possible if a device in the network sends fragmented frames, which appears uncommon in practice.

Several implementation flaws have no preconditions and are trivial to exploit. These, among others, are highlighted in bold and enable an adversary to inject network packets. Also of special interest is implementation vulnerability CVE-2020-26139 in APs since, in combination with the A-MSDU design flaw (24588) in clients, this can be abused to inject network packets without having to fulfill any special preconditions.

Packet Injection: Malicious DNS Server

If network packets can be injected towards a client, this can be abused to trick the client into using a malicious DNS server:

IPv6 Network If the target network supports IPv6, the adversary can inject ICMPv6 Router Advertisements towards the client that contain a malicious DNS server. When the client starts using this malicious DNS server, the adversary can intercept all IP-based traffic (both IPv4 and IPv6).

To spoof the ICMPv6 Router Advertisement against certain clients, the adversary has to include the IPv6 prefix of the network. Preliminary tests show that these addresses can be learned by injecting a multicast ICMPv6 ping request with as destination address `ff02::1` and with as source address the adversary's server. This causes all clients to send replies to the adversary's server, thereby revealing their IPv6 addresses.

IPv4 Network If the target network supports IPv4, the adversary can likely spoof DHCPv4 packets towards the client. To accomplish this, the adversary must predict the DHCP Transaction ID (XID) in the client's request packets. Preliminary tests indicate that several clients use predictable XIDs, e.g., generated using libc's `rand` function with as seed the system's uptime. As a result, an adversary can predict the client's XID values and spoof DHCP replies. The attacker can then include a malicious DNS server in the spoofed DHCP replies, in turn allowing the attacker to intercept IPv4 traffic.

To spoof the DHCP replies, the adversary also needs to know the IPv4 addresses being used by the target network. These addresses can be learned by injecting a guess for every possible address being used, and detecting a correct guess by the length of the encrypted replies.

Against certain devices with an IPv4/6 dual-stack, such as Linux, iOS, and macOS, it is possible to spoof ICMPv6 Router Advertisements with a DNS server address equal to an IPv4-mapped IPv6 address. For example, the router advertisement can include `::ffff:1.2.3.4` as the DNS server, causing the target to use the IPv4 address 1.2.3.4 as its DNS server.

Impact If an adversary can inject packets towards a client, you have to assume the client can be tricked into using a malicious DNS server. This can subsequently be abused to intercept and modify all IPv4 or IPv6 traffic.

Packet Injection: NAT Hole Punching

If network packets can be injected towards an AP, the adversary can abuse this to bypass the NAT/firewall and directly connect to any device in the local network. This means that if there is a sensitive local service listening on a UDP or TCP port behind the NAT (e.g. an IoT device, security camera, network storage, etc) the adversary can directly communicate with it. There are two known techniques to accomplish this:

Crafted TCP SYNs An attacker can inject a TCP SYN with the source address of the sensitive service and the destination of an attacker controlled endpoint on the internet to establish a NAT/firewall entry. Preliminary results indicate that against most NAT implementations it is then possible for the attacker to reach the service from their controlled endpoint. If the NAT implementation performs extra checks to validate the TCP state machine, additional packets can possibly be injected by the Wi-Fi attacker, or sent from the controlled endpoint, to satisfy the NAT/firewall's validation requirements and establish a connection.

Application Level Gateways Protocols such as FTP and SIP require special care when used behind a NAT, since they only work when the NAT allows incoming connections from the remote endpoint. To assure that the NAT routes the incoming remote connections to the correct local device, the NAT monitors the control traffic of the protocol and dynamically creates port mappings to track and allow incoming connections. This is often called an Application Level Gateway (ALG), and can be abused by injecting control packets to an attacker controlled endpoint which as source address the local service we want to connect to. The NAT/firewall will then create a mapping between the attacker's endpoint and any port on the targeted local device. Preliminary experiments confirmed that all routers support this functionality. This enables the adversary to directly connect to any sensitive local service listening on a UDP or TCP port.

Impact The above attacks show that packet injection can be abused to bypass the NAT/firewall and directly connect with devices in the local network. This allows an attacker to launch attacks against insecure devices behind a NAT/firewall, as long as the attacker controls a compromised device within Wi-Fi range of the target device. Using this it is possible to create worms that attack vulnerable devices that would previously be unreachable behind a NAT/firewall. The compromised devices can subsequently attack other nearby vulnerable networks. This may have a significant impact on densely populated areas where many networks are within range of each other.