Wi-Fi Protected Access® Security Considerations

May 2021

Introduction

Part of the Wi-Fi Protected Access® (WPA™) family of technologies, WPA3™-Personal provides next generation security for private Wi-Fi® networks based on a simple password credential. WPA3 raises the bar on Wi-Fi network security. To realize the security benefits available through WPA3, it is important that the implementation guidance in this document be followed. This document covers several areas deserving special mention, as well as recommended implementation considerations.

WPA3-Personal Recommendations

- Passwords used with WPA3-Personal should be complex enough to not be easily guessable. WPA3-Personal implementations should limit authentication attempts when an implementation identifies an active attack (see Password Strength below).
- Access point (AP) implementations should handle Simultaneous Authentication of Equals (SAE) operations on non-privileged processing queues which, if overwhelmed, will not result in a failure of the entire basic service set (BSS) through central processing unit (CPU) resource consumption (see Denial of Service Protection).
- SAE Diffie-Hellman Group implementation recommendations:
  - Must use only Diffie-Hellman groups 15-21; group 19 is mandatory others are optional (see Suitable Diffie-Hellman Groups and Modular Exponential (MODP Group Timing Side-Channels).
  - Should only offer a Diffie-Hellman group whose strength estimate is greater than or equal to the encryption cipher being offered (see Table 1: Diffie-Hellman Group Suitability
    *Strength estimate is a maximum value and can be decreased based on entropy estimates (Implementation Guidance for FIPS140-2 and the Cryptographic Module Validation Program)
  - Diffie-Hellman Group Downgrade).
- SAE implementations must set the security parameter k to a value of at least forty (40) per the recommendation in RFC 7664 "Dragonfly Key Exchange" for all groups to prevent timing leaks.
- SAE implementations must avoid differences in code execution that allow side channel information collection through the cache (see Cache-Based Elliptic Curve Side-Channels).
- If WPA3-Personal Transition Mode does not meet the security requirements for a deployment, WPA3-Personal and WPA2™-Personal should be deployed on individual service set identifiers (SSIDs) using unique passwords and logically separated/isolated network segments (see WPA3-Personal Transition Mode).

Failure to implement these recommendations correctly may expose the vendor implementation to attack and/or compromise the network.

Extensible Authentication Protocol Password (EAP-pwd) Recommendations

Although EAP-pwd is not currently part of WPA3, this document covers several areas related to EAP-pwd that deserve special mention, as well as recommended implementation considerations for those who choose to implement it.

While EAP-pwd may be used in enterprise deployments where authentication happens between the supplicant and the EAP-server, the SAE implementation recommendations also apply to EAP-pwd and references to "AP" are replaced by "EAP-server".

Security Considerations Detail

Password Strength

Recommendation:

Passwords used with WPA3-Personal should be complex enough to not be easily guessable, and WPA3-Personal implementations should limit authentication attempts when an implementation identifies an active attack.

Summary:

WPA3-Personal replaces the WPA2-Personal Pre-Shared Key (PSK) authentication with SAE. Unlike PSK, SAE is resistant to offline dictionary attacks. The only way for an attacker to learn a password is through repeated active attacks, each of which tests whether a single guess of the password is correct or not. Repeated authentication failures may indicate that an active attack is underway, allowing implementations to respond appropriately, including throttling authentication attempts and/or issuing alerts such as Simple Network Management Protocol (SNMP) trap, log message, or others.

The requirement for exceedingly long, random passwords with mixed-case characters and special characters is no longer necessary with WPA3-Personal. Passwords used with WPA3-Personal should be extremely difficult to guess due to the possibility of an active attack; however, the difficulty in guessing a password directly correlates to the security that SAE offers.

To illustrate the benefits that WPA3-Personal affords, consider a password selected randomly from 5,000 possible passwords. The attacker knows this but does not know which password was randomly chosen. With WPA2-Personal an attacker could determine the password through an off-line dictionary attack with a probability of success of 1. With WPA3-personal, the attacker must launch repeated active attacks, guessing a different password each time. The probability of success of the WPA3-Personal attack would only reach 0.5 after 2,500 active attacks. It should be possible to detect such an attack on WPA3-Personal long before the probability of success becomes high.

Implementations of WPA3-Personal should limit authentication attempts for a particular password—identified with an SAE Password Identifier or not—when an active attack is identified. Determination of whether an attack is underway is implementation dependent and left up to the vendor. One possible mitigation strategy may be that the AP temporarily disable a password after a series of unsuccessful authentication attempts. Note that the source medium access control (MAC) address used with failed authentication attempts is irrelevant and should not factor into the decision to disable or limit authentication for a particular password because an attacker can easily change the MAC address between attempts.

Denial of Service Protection

Recommendation:

WPA3-Personal implementations should handle SAE operations on non-privileged processing queues which, even if overwhelmed, will not result in a failure of the entire BSS through CPU resource consumption.

Summary:

An AP performs a significant amount of cryptographic work upon receipt of the first message in an SAE handshake. A denial of service attack can be initiated by flooding the AP with fraudulent messages from fake MAC addresses resulting in the failure of the entire BSS through CPU resource consumption.

SAE defines an anti-clogging cookie response in which the AP statelessly generates a string that is bound to the sender of the message when the AP detects it is under a denial of service attack. An AP may consider itself under a denial of service attack when the number of nascent connections, those in which the first message has been received but not the third message, reaches a threshold. The AP, when in a “cookie demanding” state, will not process the first SAE message unless that message contains a valid cookie bound to the MAC address of the sender.

This technique works against rudimentary and simple packet spraying attacks because the attacker is simply sending random packets and not processing responses. However, this technique does not work if the attacker chooses to receive the AP cookie request and respond with the cookie from the same MAC address. Therefore, SAE does not afford adequate protection against more sophisticated denial of service attacks. WPA3-Personal implementations should handle SAE operations on non-privileged processing queues which, even if overwhelmed, will not result in a failure of the entire BSS through CPU resource consumption.

Suitable Diffie-Hellman Groups

Recommendation:

SAE implementations must use only Diffie-Hellman groups 15-21, and group 19 is mandatory.

SAE implementations must set the security parameter k to a value of at least forty (40) as per the recommendation in RFC 7664 "Dragonfly Key Exchange" for all groups to prevent timing leaks.

Summary:

SAE performs public key cryptography using named Diffie-Hellman groups. The IKEv1 (RFC 2409) group registry maintained by the Internet Assigned Numbers Authority (IANA) maps the group's complete domain parameter set to a reference number. Not all registered groups are suitable for use with SAE.

The rules used to evaluate the suitability of groups are:

1. No binary elliptic curve (EC2N) groups
2. No groups defined over a prime field (MODP) with a prime less than 3072 bits
3. No groups defined over a prime field (MODP) with a small sub-group of prime order
4. No elliptic curve group with a prime less than 256-bits
5. No elliptic curve group that might expose detectible timing differences when used in conjunction with the SAE.

The following table indicates the recommended groups to be used with SAE. All other groups must not be used with SAE.

Group Number	Description	Strength Estimate*	Suitability
15	3072-bit MODP group	128	Suitable
16	4096-bit MODP group	152	Suitable
17	6144-bit MODP group	176	Suitable
18	8192-bit MODP group	200	Suitable
19	256-bit random ECP group (NIST)	128	Suitable (Mandatory)
20	384-bit random ECP group (NIST)	192	Suitable
21	512-bit random ECP group (NIST)	256	Suitable

Table 1: Diffie-Hellman Group Suitability

**Strength estimate is a maximum value and can be decreased based on entropy estimates (Implementation Guidance for FIPS140-2 and the Cryptographic Module Validation Program)*

Diffie-Hellman Group Downgrade

Recommendation:

SAE implementations should only offer a Diffie-Hellman group whose strength estimate is greater than or equal to the strength estimate of the encryption cipher being offered.

Summary:

In SAE, the initiator chooses the group to use and includes the group number in the first message. The responder accepts the group or responds with a message containing an error code indicating group rejection if the responder does not want to use the group. If the group is rejected, the initiator chooses another group and tries again.

This technique opens the protocol to a downgrade attack where an attacker impersonates the AP and responds with a rejection of a stronger group until the client device offers a weak group and then lets the protocol proceed with the real AP.

SAE does not inherently protect against Diffie-Hellman Group Downgrade attacks, however they can be mitigated by not allowing weak groups and only allowing rejections to offer “upgraded” groups.

Suitable Diffie-Hellman groups for use with SAE all generate a key whose strength is appropriate for the default and mandatory-to-implement cipher, AES-CCM-128. While AES-CCM-256 and AES-GCM-256 ciphers may be used with SAE, SAE uses SHA256 for key derivation thereby mitigating, to some extent, the strength benefits afforded by different groups such as group 20 or group 21. SAE implementations should only offer a Diffie-Hellman group whose strength estimate is greater than or equal to the strength estimate of the encryption cipher being offered**.** See National Institute of Standards and Technology (NIST) Special Publication SP 800-56A, Revision 3, April 2018 Appendix D Table 24 and Table 25 for more information.

MODP Group Timing Side-Channels

Recommendation:

SAE implementations must disable the use of all MODP groups with a prime less than 3072 bits to prevent side-channel timing attacks.

Summary:

The password element generation algorithm for MODP groups is affected by timing side-channels, and the obtained information can later be used to recover the password. MODP groups 22, 23, and 24 have a small sub-group and are known to be weak; refer to "Measuring small sub-group attacks against Diffie-Hellman" by Valeta et al, 2017.

See Table 1: Diffie-Hellman Group Suitability in this document for group suitability with SAE and EAP-pwd.

Cache-Based Elliptic Curve Side-Channels

Recommendation:

SAE implementations must avoid differences in code execution that allow side channel information collection through the cache.

Two methods exist:

1. Implement SAE in such a way to use constant time operations that use the same memory access pattern regardless of the values derived from the password.
2. Reduce the visibility of side channel information, for instance, by preventing sharing of cache lines between processes if efficiently supported by the hardware architecture.

Summary:

This vulnerability requires monitoring of cache access patterns on a compromised machine, one running the attacker’s software. The obtained information can later be used to recover the password. The goal is to learn if the quadratic residue (QR) test in the first iteration of the hash to curve algorithm succeeded or not. This information can be used in the offline password partitioning attack to recover the target’s password. The implementation of the hash to curve algorithm for elliptic-curve cryptography (ECC) groups does include mitigations against side channel attacks. Those mitigations include performing extra dummy iterations on random data and blinding of the underlying cryptographic calculation of the quadratic residue test. Preventing the installation of malicious software may be an effective additional mitigation approach for some device categories.

WPA3-Personal Transition Mode

Recommendation:

If WPA3-Personal Transition Mode does not meet the security requirements for a deployment, WPA3-Personal and WPA2-Personal should be deployed on individual SSIDs using unique passwords and logically separated/isolated network segments.

Summary:

WPA3-Personal Transition Mode was defined by Wi-Fi Alliance® to minimize user disruption and provide a gradual migration path to WPA3-Personal while maintaining interoperability with WPA2-Personal only devices. Since SAE is a new Wi-Fi authentication protocol and is not backward compatible with PSK, mandating WPA3-Personal only on a BSS would require every client device to support WPA3-Personal, disrupting currently deployed WPA2-Personal only devices. Once WPA3-Personal availability reaches a sufficient level amongst client devices, network owners should disable WPA3-Personal Transition Mode to achieve the full benefits of WPA3-Personal.

WPA3-Personal Transition Mode supports both WPA3-Personal and WPA2-Personal authentication on the same BSS with the same SSID, using the same password. This was done for ease-of-use and because it is not possible to make valid assumptions about user experience across diverse device types, or the security awareness of users, that would ensure a smooth rollout. A trade-off is that the common password of a WPA3-Personal Transition Mode network can be determined by attacking a WPA2-Personal device using a simple offline dictionary attack. The WPA2-Personal attack could be performed passively on a legacy client device that only supports WPA2-Personal, or a more complex active downgrade attack could be performed on a client that supports WPA3-Personal.

The passive attack on legacy WPA2-Personal only client devices is the same as exists with legacy WPA2-Personal only networks. The active attack on an WPA3-Personal client device is complex and gains the attacker little because of the possibility to run the simpler passive attack on legacy clients. An attacker who determines the password can access the network simply by using WPA2-Personal, irrespective of WPA3-Personal. In addition, even after this attack is successful and the attacker determines the password, the clients that connect with WPA3-Personal will still benefit from the forward-secrecy that SAE affords—that is, the traffic encryption keys will still remain unknown even if the password is known. This is not an attack against WPA3-Personal.

If WPA3-Personal Transition Mode does not meet the security requirements for a deployment, WPA3-Personal and WPA2-Personal should be deployed on individual SSIDs using unique passwords and logically separated/isolated network segments; network segmentation strategies and implementations are determined by the vendor. APs must be configured to support the WPA3-Personal only network to benefit from the enhanced security of SAE.

Manipulation on A-MSDU flag in QoS header

Recommendation:

A receiving device should discard all subframes in an A-MSDU if its first subframe exhibits any of the following behaviors:

- DA does not map to 802.11 MAC header RA in a frame exiting DS (i.e., From DS subfield is 1 and To DS subfield is 0 in MAC header, DA is neither the device's RA address nor a group/multicast address)
- SA does not match 802.11 MAC header TA in a frame destined to DS (i.e., To DS subfield is 1 and From DS subfield is 0 in MAC header)
- DA is AA:AA:03:00:00:00 (any DS bits including 4-addr)

Note: This does not apply to some cases when operating as a GLK STA or S1G STA.

Summary:

By flipping the A-MSDU Present subfield in the QoS Control field of the 802.11 MAC header in a non-A-MSDU frame, one can make a vulnerable receiving device accept it as an A-MSDU frame. The first subframe will exhibit the pattern as shown in the recommendation above and is usually discarded given its invalid construct. However, if a vulnerable device retains the subsequent subframes, an adversary can inject specially formulated data to solicit sensitive user information. Thus, a device should implement the mitigation mechanism to detect an abnormal first A-MSDU subframe behavior and then discard the subsequent subframes. Such a mechanism can be accomplished on the receiving device side.

A more comprehensive solution may be to protect the A-MSDU Present subfield in the QoS Control field of the 802.11 MAC header using a mechanism such as Signaling and Payload Protection (SPP), and to offer a transition means to bridge the gap between legacy and enhanced devices. The purpose of SPP is to protect an A-MSDU against attacks that manipulate the unauthenticated A-MSDU Present subfield in its plaintext QoS Control field. SPP includes this flag as part of the AAD calculation, which can effectively detect such manipulation.

Fragments encrypted with different keys

Requirement:

A device shall not assemble fragments encrypted with different keys into the same MSDU/MMPDU.

Summary:

This vulnerability stems from different keys being used to protect MSDUs. Although it is a good security practice to ensure using one key for all fragments of a frame, it is not strictly followed in some device implementations, especially in the midst of a transitional period such as key renewal. An adversary can exploit this to inject additional fragments for attack with different keys when a key renewal happens. Thus, a receiving device shall check that only one key is used for decryption of all fragments of an MSDU/MMPDU. A receiving device may implement this requirement by discarding a fragmented MSDU/MMPDU if its fragments are encrypted with different keys.

Cache attacks on frame fragments

Requirement:

If a new association or reassociation happens, a receiving device shall discard any fragments from an incomplete MSDU/MMPDU from previous association.

Summary:

If a receiving device retains previous fragments upon a new association or reassociation, an adversary can inject malicious fragments into memory to be included with new fragments after a new association. A device can effectively stop this type of attack by clearing incomplete fragments from memory upon a new association or reassociation

Non-consecutive PN fragments

Requirement:

During defragmentation, a receiving device shall check that PN (packet number) increments by exactly 1 for consecutive fragments. If not, the receiving device shall discard any fragment that does not follow this requirement.

Summary:

If a device reassembles fragments during the defragmentation process, it shall check that fragments of an MSDU/MMPDU have PNs that are ordered with an increment of exactly 1. If a device cannot observe this order from the assembly of the fragments, then the device shall discard the unexpected fragment(s). Without stringent detection of non-consecutive packet numbered fragments, an adversary can abuse this vulnerability by adding intentional fragments in between to achieve the goal of exfiltration.

Plaintext fragments in a protected network

Requirement:

When an MSDU or MMPDU is encrypted, a receiving device shall verify that every fragment from the respective MSDU/MMPDU is encrypted. Otherwise, the device shall discard the unencrypted fragments.

Summary:

Some devices in a protected network do not differentiate plaintext fragments from that of encrypted ones for an MSDU/MMPDU during defragmentation. This offers an attacker a possible way to mix intended plaintext fragments with the normal encrypted ones. To avoid such a situation, a receiving device shall not accept plaintext fragments when the corresponding MSDU/MMPDU is expected to be encrypted.

Accepting plaintext broadcast or multicast fragments

Requirement:

Broadcast or multicast frames shall not be fragmented, and a receiving device shall discard such fragments upon reception.

Summary:

Broadcast or multicast frames shall not be fragmented. However, research has found that vulnerable devices accept plaintext broadcast fragments. This allows adversary to inject arbitrary plaintext fragments into the network during and after the 4-way handshake.

Accepting plaintext A-MSDU frames that start with an EAPOL LLC/SNAP header

Requirement:

A device shall apply the frame protection rules to all the A-MSDU subframes individually. A receiving device shall discard subframes with EtherType other than EAPOL from a plaintext A-MSDU if the network uses encryption.

Summary:

Due to the fact that initial plaintext 4-way handshake frames are accepted, an attacker can disguise a plaintext A-MSDU under a valid EAPOL LLC/SNAP header, i.e., the first 8 bytes correspond to a valid RFC1042 (EAPOL LLC/SNAP) header. A vulnerable device can process subsequent plaintext subframes, allowing an adversary to inject arbitrary plaintext data and to bypass secured state of the connection.

Plaintext frame attack in a protected network

Requirement:

A receiving device shall discard plaintext Data frames in a protected network, except for EAPOL frames during the initial 4-way handshake. A receiving device shall discard plaintext robust Management frames when PMF is enabled.

Summary:

Vulnerable devices accept plaintext Data frames and plaintext robust Management frames even though security is negotiated.

Plaintext fragmented frames in a protected network

Requirement:

When an MSDU or MMPDU is expected to be encrypted, a receiving device shall discard any unencrypted fragments.

Summary:

Vulnerable devices accept plaintext fragmented frames even though a secure connection is established. This allows an attacker to inject intended plaintext frames and ignore the protected state of the network. Therefore, a device shall not accept plaintext fragments when the corresponding MSDU/MMPDU is encrypted.

Forwarding EAPOL frames

Requirement:

An AP shall not forward EAPOL frames.

Summary:

This attack is only applicable to Access Points (APs) that forward EAPOL frames. An adversary can launch denial-of-service attacks by flooding the connected network clients with such frames. This also allows the attacker to further exploit other vulnerabilities in clients connected to the vulnerable APs.

TKIP MIC of fragmented frames

Requirement:

A receiving device shall perform TKIP MIC check and if it fails at the MSDU level, the device shall discard the MSDU and invoke countermeasures as appropriate.

Summary:

Some devices do not perform MIC(message integrity code) check for fragmented TKIP frames. Such a vulnerability can be abused by injecting frames and facilitating further attacks in a network that supports TKIP. A device shall not skip MIC verification and shall discard affected MSDU(s) upon failure and perform appropriate countermeasures.

Treating fragmented frames as full frames

Requirement:

A device shall support defragmentation of a frame according to the 802.11 specification. During defragmentation, if a device detect errors in processing any of the fragments of an MSDU or MMPDU then it shall discard the invalid fragments.

Summary:

Some devices treat fragments from a frame independently as full frames, by ignoring the More Fragments subfield in the Frame Control field of the 802.11 MAC header of a frame. In addition, some devices do not support the (de)fragmentation process. These devices can still process received fragmented frames by treating them as independent full frames. This opens up a possibility for an attacker to inject malicious data in retained fragments. Thus, devices shall implement the defragmentation correctly. They shall also keep track of all the fragments belonging to an MSDU/MMPDU.

About Wi-Fi Alliance®

www.wi-fi.org

Wi-Fi Alliance® is the worldwide network of companies that brings you Wi-Fi®. Members of our collaboration forum come together from across the Wi-Fi ecosystem with the shared vision to connect everyone and everything, everywhere, while providing the best possible user experience. Since 2000, Wi-Fi Alliance has completed more than 50,000 Wi-Fi certifications. The Wi-Fi CERTIFIED™ seal of approval designates products with proven interoperability, backward compatibility, and the highest industry-standard security protections in place. Today, Wi-Fi carries more than half of the internet's traffic in an ever-expanding variety of applications. Wi-Fi Alliance continues to drive the adoption and evolution of Wi-Fi, which billions of people rely on every day.