Internet Engineering Task Force (IETF)
Request for Comments: 9797
Category: Informational
ISSN: 2070-1721

J. Henry
Cisco Systems
Y. Lee
Comcast

June 2025

Randomized and Changing Media Access Control (MAC) Addresses: Context, Network Impacts, and Use Cases

Abstract

To limit the privacy issues created by the association between a device, its traffic, its location, and its user in IEEE 802 networks, client vendors and client OS vendors have started implementing Media Access Control (MAC) address randomization. This technology is particularly important in Wi-Fi networks (defined in IEEE 802.11) due to the over-the-air medium and device mobility. When such randomization happens, some in-network states may break, which may affect network connectivity and user experience. At the same time, devices may continue using other stable identifiers, defeating the purpose of MAC address randomization.

This document lists various network environments and a range of network services that may be affected by such randomization. This document then examines settings where the user experience may be affected by in-network state disruption. Last, this document examines some existing frameworks that maintain user privacy while preserving user quality of experience and network operation efficiency.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at https://www.rfc-editor.org/info/rfc9797.

Copyright Notice

Table of Contents

1. Introduction

When the MAC address was first introduced in [IEEE_802], it was used in wired Ethernet networks [IEEE_802.3]. Due to the nature of wired networks, devices were relatively stationary, and the physical connection imposed a boundary that restricted attackers from easily accessing the network data. However, [IEEE_802.11] (Wi-Fi) brought new challenges when it was introduced.

The flexibility of Wi-Fi technology has revolutionized communications and become the preferred, and sometimes the only, technology used by devices such as laptops, tablets, and Internet of Things (IoT) devices. Wi-Fi is an over-the-air medium that allows attackers with surveillance equipment to monitor WLAN packets and track the activity of WLAN devices. It is also sometimes possible for attackers to monitor the WLAN packets behind the Wi-Fi Access Point (AP) over the wired Ethernet. Once the association between a device and its user is made, identifying the device and its activity is sufficient to deduce information about what the user is doing, without the user's consent.

To reduce the risks of identifying a device only by the MAC address, client OS vendors have started implementing Randomized and Changing MAC addresses (RCM). By randomizing the MAC address, it becomes harder to use the MAC address to construct a persistent association between a flow of data packets and a device, assuming no other visible unique identifiers or stable patterns are in use. When individual devices are no longer easily identifiable, it also becomes difficult to associate a series of network packet flows in a prolonged period with a particular individual using one specific device if the device randomizes the MAC address governed by the OS privacy policies.

However, such address changes may affect the user experience and the efficiency of legitimate network operations. For a long time, network designers and implementers relied on the assumption that a given machine, in a network implementing IEEE 802 technologies [IEEE_802], would be represented by a unique network MAC address that would not change over time. When this assumption is broken, network communication may be disrupted. For example, sessions established between the end device and the network services may break, and packets in transit may suddenly be lost. If multiple clients implement aggressive (e.g., once an hour or shorter) MAC address randomization without coordination with network services, some network services, such as MAC address caching in the AP and the

upstream Layer 2 switch, may not be able to handle the load, which may result in an unexpected network interruption.

At the same time, some network services rely on the end station (as defined by [IEEE_802]) to provide an identifier, which can be the MAC address or another value. This document also refers to the end station as a "device" or "machine". If the client implements MAC address randomization but continues sending the same static identifier, then the association between a stable identifier and the station continues despite the RCM scheme. There may be environments where such continued association is desirable, but there may be others where user privacy has more value than any continuity of network service state.

It is useful for implementations of client and network devices to enumerate services that may be affected by RCM and to evaluate possible frameworks to maintain both the quality of user experience and network efficiency while RCM happens and user privacy is strengthened. This document presents these assessments and recommendations.

Although this document mainly discusses MAC address randomization in Wi-Fi networks [IEEE_802.11], the same principles can be easily extended to any IEEE 802 networks [IEEE_802].

This document is organized as follows:

* Section 2 discusses the current status of using MAC address as identity.

* Section 3 discusses various actors in the network that will be impacted by MAC address randomization.

* Section 4 examines the degrees of trust between personal devices and the entities at play in a network domain.

* Section 5 discusses various network environments that will be impacted.

* Section 6 analyzes some existing network services that will be impacted.

* Appendix A includes some existing frameworks.

2. MAC Address as Identity: User vs. Device

In IEEE 802 [IEEE_802] technologies, the Media Access Control (MAC) layer defines rules to control how a device accesses the shared medium. In a network where a machine can communicate with one or more other machines, one such rule is that each machine needs to be identified as either the target destination of a message or the source of a message (and the target destination of the answer). Initially intended as a 48-bit (6-octet) value in the first versions of IEEE 802, other standards under the IEEE 802 [IEEE_802] umbrella allow this address to take an extended format of 64 bits (8 octets), which enabled a larger number of MAC addresses to coexist as IEEE 802 technologies became widely adopted.

Regardless of the address length, different networks have different needs, and several bits of the first octet are reserved for specific purposes. In particular, the first bit is used to identify the destination address as an individual (bit set to 0) or a group address (bit set to 1). The second bit, called the Universal/Local (U/L) address bit, indicates whether the address has been assigned by a universal or local administrator. Universally administered addresses have this bit set to 0. If this bit is set to 1, the

entire address is considered to be locally administered (see Clause 8.4 of [IEEE_802]). Note that universally administered MAC addresses are required to be registered with the IEEE, while locally administered MAC addresses are not.

The intent of this provision is important for the present document. [IEEE_802] recognizes that some devices (e.g., smart thermostats) may never change their attachment network and will not need a globally unique MAC address to prevent address collision against any other device in any other network. The U/L bit can be set to signal to the network that the MAC address is intended to be locally unique (not globally unique). [IEEE_802] did not initially define the MAC address allocation schema when the U/L bit is set to 1. It states the address must be unique in a given broadcast domain (i.e., the space where the MAC addresses of devices are visible to one another).

It is also important to note that the purpose of the universal version of the address was to avoid collisions and confusion, as any machine could connect to any network, and each machine needs to determine if it is the intended destination of a message or its response. Clause 8.4 of [IEEE_802] reminds network designers and operators that all potential members of a network need to have a unique identifier in that network (if they are going to coexist in the network without confusion on which machine is the source or destination of any message). The advantage of an administrated address is that a node with such an address can be attached to any Local Area Network (LAN) in the world with an assurance that its address is unique in that network.

With the rapid development of wireless technologies and mobile devices, this scenario became very common. With a vast majority of networks implementing IEEE 802 radio technologies [IEEE_802] at the access, the MAC address of a wireless device can appear anywhere on the planet and collisions should still be avoided. However, the same evolution brought the distinction between two types of devices that [IEEE_802] generally refers to as "nodes in a network" (see Section 6.2 of [IEEE_802E] for definitions of these devices):

Shared Service Device: A device used by enough people that the device itself, its functions, or its traffic cannot be associated with a single or small group of people. Examples of such devices include switches in a dense network, (WLAN) access points [IEEE_802.11] in a crowded airport, and task-specific devices (e.g., barcode scanners).

Personal Device: A machine or node primarily used by a single person or small group of people, so that any identification of the device or its traffic can also be associated with the identification of the primary user or their online activity.

Identifying the device is trivial if it has a unique MAC address. Once this unique MAC address is established, detecting any elements that directly or indirectly identify the user of the device (i.e., Personally Identifiable Information (PII)) is enough to link the MAC address to that user. Then, any detection of traffic that can be associated with the device will also be linked to the known user of that device (i.e., Personally Correlated Information (PCI)).

2.1. Privacy of MAC Addresses

The possible identification or association presents a privacy issue, especially with wireless technologies. For most of them ([IEEE_802.11] in particular), the source and destination MAC addresses are not encrypted even in networks that implement encryption. This lack of encryption allows each machine to easily detect if it is the intended target of the message before attempting

to decrypt its content and also helps identify the transmitter in order to use the right decryption key when multiple unicast keys are in effect.

This identification of the user associated with a node was clearly not the intent of the IEEE 802 MAC address. A logical solution to remove this association is to use a locally administered address instead and change the address in a fashion that prevents a continuous association between one MAC address and some PII. However, other network devices on the same LAN implementing a MAC layer also expect each device to be associated with a MAC address that would persist over time. When a device changes its MAC address, other devices on the same LAN may fail to recognize that the same machine is attempting to communicate with them. This type of MAC address is referred to as 'persistent' MAC address in this document. This assumption sometimes adds to the PII confusion, for example, in the case of Authentication, Authorization, and Accounting (AAA) services [RFC3539] authenticating the user of a machine and associating the authenticated user to the device MAC address. Other services solely focus on the machine (e.g., DHCPv4 [RFC2131]) but still expect each device to use a persistent MAC address, for example, to reassign the same IP address to a returning device. Changing the MAC address may disrupt these services.

3. The Actors: Network Functional Entities and Human Entities

The risk of service disruption is weighed against the privacy benefits. However, the plurality of actors involved in the exchanges tends to blur the boundaries of which privacy violations should be protected against. Therefore, it is useful to list the actors associated with the network exchanges because they either actively participate in these exchanges or can observe them. Some actors are functional entities, while others are human (or related) entities.

3.1. Network Functional Entities

Network communications based on IEEE 802 technologies commonly rely on station identifiers based on a MAC address. This MAC address is utilized by several types of network functional entities such as applications or devices that provide a service related to network operations.

1. Wireless access network infrastructure devices (e.g., WLAN access points or controllers): These devices participate in IEEE 802 LAN operations. As such, they need to identify each machine as a source or destination to successfully continue exchanging frames. As a device changes its network attachment (roams) from one access point to another, the access points can exchange contextual information (e.g., device MAC address and keying material), allowing the device session to continue seamlessly. These access points can also inform devices further in the wired network about the roam to ensure that Layer 2 frames are redirected to the new device access point.

2. Other network devices operating at the MAC layer: Many wireless network access devices (e.g., access points [IEEE_802.11]) are conceived as Layer 2 devices, and as such, they bridge a frame from one medium (e.g., Wi-Fi [IEEE_802.11]) to another (e.g., Ethernet [IEEE_802.3]). This means that the MAC address of a wireless device often exists on the wire beyond the wireless access device. Devices connected to this wire also implement IEEE 802.3 technologies [IEEE_802.3], and as such, they operate on the expectation that each device is associated with a MAC address that persists for the duration of continuous exchanges. For example, switches and bridges associate MAC addresses to individual ports (so as to know to which port to send a frame

intended for a particular MAC address). Similarly, AAA services can validate the identity of a device and use the device MAC address as the first pointer to the device identity (before operating further verification). Similarly, some networking devices offer Layer 2 filtering policies that may rely on the connected MAC addresses. IEEE 802.1X-enabled devices [IEEE_802.1X] may also selectively put the interface in a blocking state until a connecting device is authenticated. These services then use the MAC address as the first pointer to the device identity to allow or block data traffic. This list is not exhaustive. Multiple services are defined for Ethernet networks [IEEE_802.3], and multiple services defined by the IEEE 802.1 working group are also applicable to Ethernet networks [IEEE_802.3]. Wireless access points may also connect using other mediums (e.g., the Data-Over-Cable Service Interface Specification (DOCSIS) [DOCSIS]) that implement mechanisms under the umbrella of the general 802 Standard and therefore expect the unique and persistent association of a MAC address to a device.

3. Network devices operating at upper layers: Some network devices provide functions and services above the MAC layer. Some of them also operate a MAC layer function. For example, routers provide IP forwarding services but rely on the device MAC address to create the appropriate frame structure. Other devices and services operate at upper layers but also rely upon the IEEE 802 principles of unique MAC-to-device mapping. For example, the Address Resolution Protocol (ARP) [RFC826] and Neighbor Discovery Protocol (NDP) [RFC4861] use a MAC address to create the mapping of an IP address to a MAC address for packet forwarding. If a device changes its MAC address without a mechanism to notify the Layer 2 switch it is connected to or is the provider of a service that expects a stable MAC-to-device mapping, the provider of the service and traffic forwarding may be disrupted.

3.2. Human-Related Entities

Humans may actively participate in the network structure and operations or be observers at any point of the network lifecycle. Humans could be users of wireless devices or people operating wireless networks.

1. Over-the-Air (OTA) observers: The transmitting or receiving MAC address is usually not encrypted in wireless exchanges using IEEE 802 technologies, and any protocol-compatible device in range of the signal can read the frame header. As such, OTA observers are able to read the MAC addresses of individual transmissions. Some wireless technologies also support techniques to establish distances or positions, allowing the observer, in some cases, to uniquely associate the MAC address with a physical device and its associated location. An OTA observer may have a legitimate reason to monitor a particular device, for example, for IT support operations. However, another actor might also monitor the same device to obtain PII or PCI.

2. Wireless access network operators: Some wireless access networks host devices that meet specific requirements, such as device type (e.g., IoT-only networks and factory operational networks). Therefore, operators can attempt to identify the devices (or the users) connecting to the networks under their care. They often use the MAC address to represent an identified device.

3. Network access providers: Wireless access networks are often considered beyond the first two layers of the OSI model. For example, a law enforcement agency (e.g., the FBI in the United States) may legally require the network access provider to identify communications from a subject. In this context, the

operating access networks need to identify the devices used by the subjects and cross-reference the data generated by the devices in the network. In other contexts, the operating access networks assign resources based on contractual conditions (e.g., fee and bandwidth fair share). In these scenarios, the operators may use the MAC address to identify the devices and the users of their networks.

4. Over-the-Wired internal (OTWi) observers: Because the device wireless MAC address continues to be present over the wire if the infrastructure connection device (e.g., access point) functions as a Layer 2 bridge, observers may be positioned over the wire and may read transmission MAC addresses. Such capability supposes that the observer has access to the wired segment of the broadcast domain where the frames are exchanged. A broadcast domain is a logical segment of a network in which devices can send, receive, and monitor data frames from all other devices within the same segment. In most networks, such capability requires physical access to an infrastructure wired device in the broadcast domain (e.g., switch closet) and is therefore not accessible to all.

5. Over-the-Wired external (OTWe) observers: Beyond the broadcast domain, frame headers are removed by a routing device, and a new Layer 2 header is added before the frame is transmitted to the next segment. The device MAC address is not visible anymore unless a mechanism copies the MAC address into a field that can be read while the packet travels to the next segment (e.g., IPv6 addresses built from the MAC address prior to the use of the methods defined in [RFC8981] and [RFC7217]). Therefore, unless this last condition exists, OTWe observers are not able to see the device MAC address.

4. Degrees of Trust

The surface of PII exposures that can drive MAC address randomization depends on (1) the environment where the device operates, (2) the presence and nature of other devices in the environment, and (3) the type of network the device is communicating through. Consequently, a device can use an identifier (such as a MAC address) that can persist over time if trust with the environment is established, or it can use an identifier that is temporary if an identifier is required for a service in an environment where trust has not been established. Note that trust is not binary. It is useful to distinguish what trust a personal device may establish with the different entities at play in a network domain where a MAC address may be visible:

1. Full trust: The device establishes a trust relationship and shares its persistent MAC address with the access network devices (e.g., access point and WLAN controller). The network provides necessary security measures to prevent observers or network actors from accessing PII. The device (or its user) also has confidence that its MAC address is not shared beyond the Layer 2 broadcast domain boundary.

2. Selective trust: Depending on the predefined privacy policies, a device may decide to use one pseudo-persistent MAC address for a set of network elements and another pseudo-persistent MAC address for another set of network elements. Examples of privacy policies can be a combination of Service Set Identifier (SSID) and Basic Service Set Identifier (BSSID), a particular time of day, or a preset time duration.

3. Zero trust: A device may randomize its MAC address with any local entity reachable through the AP. It may generate a temporary MAC address to each of them. That temporary MAC address may or may

not be the same for different services.

5. Environments

The trust relationship depends on the relationship between the user of a personal device and the operator of a network service that the personal device may use. It is useful to observe the typical trust structure of common environments:

(A) Residential settings under the control of the user: This is a typical home network with Wi-Fi in the LAN and Internet in the WAN. In this environment, traffic over the Internet does not expose the MAC address of the internal device if it is not copied to another field before routing happens. The wire segment within the broadcast domain is under the control of the user and is usually not at risk of hosting an eavesdropper. Full trust is typically established at this level among users and with the network elements. Note that "Full trust" in this context is referring to the MAC address persistency. It does not extend to full trust between applications or devices. The device trusts the access point and all Layer 2 domain entities beyond the access point, where the Wi-Fi transmissions can be detected, but there is no guarantee that an eavesdropper will not observe the communications. As such, even in this environment, it is common to assume that attackers may still be able to monitor unencrypted information such as MAC addresses. If a device decides to not fully trust the network, it might apply any necessary policy to protect its identity. Most users connecting to a residential network only expect simple Internet connectivity services, so the network services are simple. If users have issues connecting to the network or accessing the Internet, they expect limited to no technical support.

(B) Managed residential settings: Examples of this type of environment include shared living facilities and other collective environments where an operator manages the network for the residents. The OTA exposure is similar to (A). The operator may be requested to provide IT support to the residents and may need to identify device activity in real time or analyze logs. The infrastructure is shared and covers a larger area than in (A); residents may connect to the network from different locations. For example, they may regularly connect to the network from their own apartments and occasionally connect from common areas. The device may decide to use different pseudo-persistent MAC addresses as described in Section 4. As such, the degree of trust is "Selective trust". In this environment, the network services will be slightly more complex than in (A). The network may be segmented by locations and multiple SSIDs. Users' devices should be able to join the network without pre-certification or pre-approval. In most cases, users only need simple connectivity; thus, network support will be slightly (but not significantly) more complicated than in (A).

(C) Public guest networks: Public hotspots in shopping malls, hotels, stores, train stations, and airports are typical examples of this environment. In this environment, trust is commonly not established with any element of the Layer 2 broadcast domain. Users do not anticipate a public guest network using the MAC address information to identify their location and network activity. They do not trust the network and do not want the network to memorize them permanently. The degree of trust is "Zero trust". Devices in this network should avoid using a long-lived MAC address to prevent fingerprinting. For example, the device may use a different MAC address every time it attaches to a new Wi-Fi access point. Some guest network operators may legally abide to identify devices. They

should not use the MAC address for such a function. Most users connecting to a public guest network only expect simple Internet connectivity services, so the network services are simple. If users have issues connecting to the network or accessing the Internet, they expect limited to no technical support. Thus, the network support level is low.

(D) Enterprises with Bring-Your-Own-Device (BYOD): This type of network is similar to (B) except that the onboarding devices are subjected to pre-approval and pre-certification. The devices are usually personal devices and are not under the control of the corporate IT team. Compared to residential networks, enterprise networks usually provide more sophisticated network services including, but not limited to, application-based and identity-based network policies. Changing the MAC address may interrupt network services if the services are based on that MAC address. Thus, network operations will be more complex, so the network support level is high.

(E) Managed enterprises: This type of network is similar to (D). The main difference is that the devices are owned and managed by the enterprise. Because both the network and the devices are owned and managed by the enterprise, the degree of trust is "Full trust". Network services and the network support level are the same as in (D).

Table 1 summarizes the environments described above.

Use Cases	Degree of Trust	Network Admin	Network Services	Network Support Expectation
(A) Residential settings under the control of the user	Full trust	User	Simple	Low
(B) Managed residential settings	Selective trust	IT	Medium	Medium
(C) Public guest networks	Zero trust	ISP	Simple	Low
(D) Enterprises with Bring-Your-Own-Device (BYOD)	Selective trust	IT	Complex	High
(E) Managed enterprises	Full trust	IT	Complex	High

Table 1: Use Cases

Existing technical frameworks that address some of the requirements of the use cases listed above are discussed in Appendix A.

6. Network Services

Different network environments provide different levels of network services, from simple to complex. At its simplest level, a network can provide a wireless connecting device with basic IP communication service (e.g., DHCPv4 [RFC2131] or Stateless Address Autoconfiguration (SLAAC) [RFC4862]) and an ability to connect to the Internet (e.g., DNS service or relay and routing in and out through a local gateway). The network can also offer more advanced services, such as managed instant messaging service, file storage, printing,

and/or local web service. Larger and more complex networks can also incorporate more advanced services, from AAA to Augmented Reality (AR) or Virtual Reality (VR) applications. To the network, its top priority is to provide the best quality of experience to its users. Often the network contains policies that help to make a forwarding decision based on the network conditions, the device, and the user identity associated to the device. For example, in a hospital private network, the network may contain a policy to give highest priority to doctors' Voice-Over-IP packets. In another example, an enterprise network may contain a policy to allow applications from a group of authenticated devices to use Explicit Congestion Notification (ECN) [RFC3168] for congestion and/or Differentiated Services Code Point (DSCP) [RFC8837] for classification to signal the network for a specific network policy. In this configuration, the network is required to associate the data packets to an identity to validate the legitimacy of the marking. Before RCM, many network systems used a MAC address as a persistent identity to create an association between user and device. After implementing RCM, the association is broken.

6.1. Device Identification and Associated Problems

Wireless access points and controllers use the MAC address to validate the device connection context, including protocol capabilities, confirmation that authentication was completed, quality of service or security profiles, and encryption keying material. Some advanced access points and controllers also include upper layer functions whose purpose is covered below. A device changing its MAC address, without another recorded device identity, would cause the access point and the controller to lose the relation between a connection context and the corresponding device. As such, the Layer 2 infrastructure does not know that the device (with its new MAC address) is authorized to communicate through the network. The encryption keying material is not identified anymore (causing the access point to fail to decrypt the device packets and fail to select the right key to send encrypted packets to the device). In short, the entire context needs to be rebuilt, and a new session restarted. The time consumed by this procedure breaks any flow that needs continuity or short delay between packets on the device (e.g., real-time audio, video, AR/VR, etc.). For example, [IEEE_802.11i] recognizes that a device may leave and rejoin the network after a short time window. As such, the standard suggests that the infrastructure should keep the context for a device for a while after the device was last seen. The device should maintain the same MAC address in such a scenario.

Some network equipment such as multi-layer routers and Wi-Fi access points, which serve both Layer 2 and Layer 3 in the same device, rely on ARP [RFC826] and NDP [RFC4861] to build the MAC-to-IP table for packet forwarding. The size of the MAC address cache in the Layer 2 switch is finite. If new entries are created faster than the old entries are flushed by the idle timer, it is possible to cause an unintentional denial-of-service attack. For example, the default timeout of the MAC address cache in Linux is set to 300 seconds. Aggressive MAC randomization from many devices in a short time interval (e.g., less than 300 seconds) may cause the Layer 2 switch to exhaust its resources, holding in memory traffic for a device whose entry can no longer be found. For the RCM device, these effects translate into session discontinuity and disrupt the active sessions. The discontinuity impact may vary. Real-time applications such as video conference may experience short interruption while non-real-time applications such as video streaming may experience minimal or no impact. The device should carefully balance when to change the MAC address after analyzing the nature of the running applications and its privacy policy.

In wireless contexts, IEEE 802.1X authenticators [IEEE_802.1X] rely on the device and user identity validation provided by a AAA server to change the interface from a blocking state to a forwarding state. The MAC address is used to verify that the device is in the authorized list and to retrieve the associated key used to decrypt the device traffic. A change in MAC address causes the port to be closed to the device data traffic until the AAA server confirms the validity of the new MAC address. Consequently, MAC address randomization can disrupt the device traffic and strain the AAA server.

Without a unique identification of the device, DHCPv4 servers [RFC2131] lose track of which IP address is validly assigned. Unless the RCM device releases the IP address before changing its MAC address, DHCPv4 servers are at risk of scope exhaustion, causing new devices (and RCM devices) to fail to obtain a new IP address. Even if the RCM device releases the IP address before changing the MAC address, the DHCPv4 server typically holds the released IP address for a certain duration, in case the leaving MAC returns. As the DHCPv4 server [RFC2131] cannot know if the release is due to a temporary disconnection or a MAC randomization, the risk of scope address exhaustion exists even in cases where the IP address is released.

Network devices with self-assigned IPv6 addresses (e.g., with SLAAC [RFC4862]) and devices using static IP addresses rely on mechanisms like Optimistic Duplicate Address Detection (DAD) [RFC4429] and NDP [RFC4861] for peer devices to establish the association between the target IP address and a MAC address, and these peers may cache this association in memory. Changing the MAC address, even at the disconnection-reconnection phase, without changing the IP address may disrupt the stability of these mappings for these peers if the change occurs within the caching period. Note that this behavior is against standard operation and existing privacy recommendations. Implementations must avoid changing the MAC address while maintaining the previously assigned IP address without consulting the network.

Routers keep track of which MAC address is on which interface so that they can form the proper Data Link header when forwarding a packet to a segment where MAC addresses are used. MAC address randomization can cause MAC address cache exhaustion but also the need for frequent Address Resolution Protocol (ARP) [RFC826], Reverse Address Resolution Protocol (RARP) [RFC903], and Neighbor Solicitation and Neighbor Advertisement [RFC4861] exchanges.

In residential settings (environment type A in Section 5), policies can be in place to control the traffic of some devices (e.g., parental control or blocklist filters). These policies are often based on the device MAC address. MAC address randomization removes the possibility for such control.

In residential settings (environment type A) and in enterprises (environment types D and E), device recognition and ranging may be used for IoT-related functionalities (e.g., door unlock, preferred light and temperature configuration, etc.) These functions often rely on the detection of the device wireless MAC address. MAC address randomization breaks the services based on such models.

In managed residential settings (environment type B) and in enterprises (environment types D and E), the network operator is often requested to provide IT support. With MAC address randomization, real-time support is only possible if the user can provide the current MAC address. Service improvement support is not possible if the MAC address that the device had at the time of the reported issue (in the past) is not known at the time the issue is reported.

In managed enterprise environments, policies are associated with each group of objects, including IoT devices. MAC address randomization may prevent an IoT device from being identified properly and thus lead to network quarantine and disruption of operations.

7. IANA Considerations

This document has no IANA actions.

8. Security Considerations

Privacy considerations are discussed throughout this document.

9. Informative References

[DOCSIS] CableLabs, "Cable Modem Operations Support System Interface Specification", Data-Over-Cable Service Interface Specifications, DOCSIS 4.0, Version I06, March 2022, <https://www.cablelabs.com/specifications/CM-SP-CM-OSSIv4.0?v=I06>.

[IEEE_802] IEEE, "IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture", IEEE Std 802-2014, DOI 10.1109/IEEESTD.2014.6847097, 30 June 2014, <https://ieeexplore.ieee.org/document/6847097>.

[IEEE_802.11] IEEE, "IEEE Standard for Information Technology--Telecommunications and Information Exchange between Systems - Local and Metropolitan Area Networks--Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", IEEE Std 802.11-2020, DOI 10.1109/IEEESTD.2021.9363693, 26 February 2021, <https://ieeexplore.ieee.org/document/9363693>.

[IEEE_802.11bh] IEEE, "IEEE Standard for Information Technology--Telecommunications and Information Exchange Between Systems Local and Metropolitan Area Networks--Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 1: Operation with Randomized and Changing MAC Addresses", IEEE Std 802.11bh-2024, DOI 10.1109/IEEESTD.2025.11023005, 3 June 2025, <https://ieeexplore.ieee.org/document/11023005>.

[IEEE_802.11i] IEEE, "IEEE 802.11i-2004 - Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: Amendment 6: Medium Access Control (MAC) Security Enhancements", IEEE Std 802.11i-2004, DOI 10.1109/10.1109/IEEESTD.2004.94585, 24 July 2004, <https://ieeexplore.ieee.org/document/1318903>.

[IEEE_802.1X] IEEE, "IEEE Standard for Local and Metropolitan Area Networks--Port-Based Network Access Control", IEEE Std 802.1X-2020, DOI 10.1109/IEEESTD.2020.9018454, 28 February 2020, <https://ieeexplore.ieee.org/document/9018454>.

[IEEE_802.3] IEEE, "IEEE Standard for Ethernet", IEEE Std 802.3-2022, DOI 10.1109/IEEESTD.2022.9844436, 29 July 2022, <https://ieeexplore.ieee.org/document/9844436>.

[IEEE_802E]
IEEE, "IEEE Recommended Practice for Privacy Considerations for IEEE 802(R) Technologies", IEEE Std 802E-2020, DOI 10.1109/IEEESTD.2020.9018454, 13 November 2020, <https://ieeexplore.ieee.org/document/9257130>.

[RADIUS] DeKok, A., "Deprecating Insecure Practices in RADIUS", Work in Progress, Internet-Draft, draft-ietf-radext-deprecating-radius-06, 25 May 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-radext-deprecating-radius-06>.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <https://www.rfc-editor.org/info/rfc2131>.

[RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <https://www.rfc-editor.org/info/rfc3168>.

[RFC3539] Aboba, B. and J. Wood, "Authentication, Authorization and Accounting (AAA) Transport Profile", RFC 3539, DOI 10.17487/RFC3539, June 2003, <https://www.rfc-editor.org/info/rfc3539>.

[RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, DOI 10.17487/RFC4429, April 2006, <https://www.rfc-editor.org/info/rfc4429>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/info/rfc4861>.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <https://www.rfc-editor.org/info/rfc4862>.

[RFC6614] Winter, S., McCauley, M., Venaas, S., and K. Wierenga, "Transport Layer Security (TLS) Encryption for RADIUS", RFC 6614, DOI 10.17487/RFC6614, May 2012, <https://www.rfc-editor.org/info/rfc6614>.

[RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <https://www.rfc-editor.org/info/rfc7217>.

[RFC826] Plummer, D., "An Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", STD 37, RFC 826, DOI 10.17487/RFC0826, November 1982, <https://www.rfc-editor.org/info/rfc826>.

[RFC8837] Jones, P., Dhesikan, S., Jennings, C., and D. Druta, "Differentiated Services Code Point (DSCP) Packet Markings for WebRTC QoS", RFC 8837, DOI 10.17487/RFC8837, January 2021, <https://www.rfc-editor.org/info/rfc8837>.

[RFC8981] Gont, F., Krishnan, S., Narten, T., and R. Draves, "Temporary Address Extensions for Stateless Address Autoconfiguration in IPv6", RFC 8981,

DOI 10.17487/RFC8981, February 2021,
<https://www.rfc-editor.org/info/rfc8981>.

[RFC903] Finlayson, R., Mann, T., Mogul, J., and M. Theimer, "A Reverse Address Resolution Protocol", STD 38, RFC 903, DOI 10.17487/RFC0903, June 1984, <https://www.rfc-editor.org/info/rfc903>.

[WBA-OPENROAMING]
Tomas, B., Grayson, M., Canpolat, N., Cockrell, B. A., and S. Gundavelli, "WBA OpenRoaming Wireless Federation", Work in Progress, Internet-Draft, draft-tomas-openroaming-05, 15 April 2025, <https://datatracker.ietf.org/doc/html/draft-tomas-openroaming-05>.

Appendix A. Existing Frameworks

A.1. IEEE 802.1X with WPA2 / WPA3

In a typical enterprise Wi-Fi environment, IEEE 802.1X authentication [IEEE_802.1X] coupled with WPA2 or WPA3 [IEEE_802.11i] encryption schemes are commonly used for onboarding a Wi-Fi device. This allows the mutual identification of the client device or the user of the device and an authentication authority. The authentication exchange does not occur in clear text, and the user or device identity can be concealed from unauthorized observers. However, in most cases, the authentication authority is under the control of the same entity as the network access provider. This may lead to exposing the user or device identity to the network owner.

This scheme is well-adapted to an enterprise environment, where a level of trust is established between the user and the enterprise network operator. In this scheme, MAC address randomization can occur through brief disconnections and reconnections (under the rules of [IEEE_802.11bh]). Authentication may then need to reoccur, with an associated cost of service disruption, an additional load on the enterprise infrastructure, and an associated benefit of limiting the exposure of a continuous MAC address to external observers. The adoption of this scheme is limited outside of the enterprise environment by the requirement to install an authentication profile on the end device, which would be recognized and accepted by a local authentication authority and its authentication server. Such a server is uncommon in a home environment, and the procedure to install a profile is cumbersome for most untrained users. The likelihood that a user or device profile would match a profile recognized by a public Wi-Fi authentication authority is also fairly limited. This may restrict the adoption of this scheme for public Wi-Fi as well. Similar limitations are found in the hospitality environment. The hospitality environment refers to space provided by the hospitality industry, which includes but is not limited to hotels, stadiums, restaurants, concert halls, and hospitals.

A.2. OpenRoaming

In order to alleviate some of the limitations listed above, the Wireless Broadband Alliance (WBA) OpenRoaming standard introduces an intermediate trusted relay between local venues (places where some public Wi-Fi is available) and sources of identity [WBA-OPENROAMING]. The federation structure extends the type of authorities that can be used as identity sources (compared to the typical enterprise-based IEEE 802.1X scheme for Wi-Fi [IEEE_802.1X]) and facilitates the establishment of trust between local networks and an identity provider. Such a procedure increases the likelihood that one or more identity profiles for the user or the device will be recognized by a local network. At the same time, authentication does not occur to the local network. This may offer the possibility for the user or

the device to keep their identity obfuscated from the local network operator, unless that operator specifically expresses the requirement to disclose such identity (in which case the user has the option to accept or decline the connection and associated identity exposure).

The OpenRoaming scheme seems well-adapted to public Wi-Fi and hospitality environments. It defines a framework to protect the identity from unauthorized entities while permitting mutual authentication between the device or the user and a trusted identity provider. Just like the standard IEEE 802.1X scheme for Wi-Fi [IEEE_802.1X], authentication allows for the establishment of WPA2 or WPA3 keys [IEEE_802.11i] that can then be used to encrypt the communication between the device and the access point. The encryption adds extra protection to prevent the network traffic from being eavesdropped.

MAC address randomization can occur through brief disconnections and reconnections (under the rules of [IEEE_802.11bh]). Authentication may then need to reoccur, with an associated cost of service disruption, an additional load on the venue and identity provider infrastructure, and an associated benefit of limiting the exposure of a continuous MAC address to external observers. Limitations of this scheme include the requirement to first install one or more profiles on the client device. This scheme also requires the local network to support RADSEC [RFC6614] and the relay function, which may not be common in small hotspot networks and home environments.

It is worth noting that, as part of collaborations between the IETF MADINAS Working Group and WBA around OpenRoaming, some RADIUS privacy enhancements have been proposed in the IETF RADEXT Working Group. For instance, [RADIUS] describes good practices in the use of Chargeable-User-Identity (CUI) between different visited networks, making it better suited for public Wi-Fi and hospitality use cases.

A.3. Proprietary RCM Schemes

Most client OS vendors offer RCM schemes that are enabled by default (or easy to enable) on client devices. With these schemes, the device changes its MAC address, when not associated, after having used a given MAC address for a semi-random duration window. These schemes also allow for the device to manifest a different MAC address in different SSIDs.

Such a randomization scheme enables the device to limit the duration of exposure of a single MAC address to observers. In [IEEE_802.11bh], MAC address randomization is not allowed during a given association session, and MAC address randomization can only occur through disconnection and reconnection. Authentication may then need to reoccur, with an associated cost of service disruption and additional load on the venue and identity provider infrastructure, directly proportional to the frequency of the randomization. The scheme is also not intended to protect from the exposure of other identifiers to the venue network (e.g., DHCP option 012 [host name] visible to the network between the AP and the DHCPv4 server).

Authors' Addresses

Jerome Henry
Cisco Systems
United States of America
Email: jerhenry@cisco.com

Yiu L. Lee
Comcast

1800 Arch Street
Philadelphia, PA 19103
United States of America
Email: yiu_lee@comcast.com