

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 27 April 2022

J. Henry
Cisco Systems
Y. Lee
Comcast
24 October 2021

Randomized and Changing MAC Address Use Cases
draft-henry-madinas-framework-03

Abstract

To limit the association between a device traffic and its user, client vendors have started implementing MAC address rotation. When such rotation happens, some in-network states may break, which may affect network efficiency and the user experience. At the same time, devices may continue sending other stable identifiers, defeating the MAC rotation purposes. This document lists various network environments and a set of network services that may be affected by such rotation. This document then examines settings where the user experience may be affected by in-network state disruption, and settings where other machine identifiers may expose the user privacy. Last, this document examines solutions to maintain user privacy while preserving user quality of experience and network operation efficiency.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 27 April 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Requirements Language	3
2.	MAC Address as an Identity: User vs. Device	3
3.	The Actors: Network Functional Entities and Human Entities .	6
3.1.	Network Functional Entities	6
3.2.	Human-related Entities	7
3.3.	The Trust and the Environments	8
3.4.	The Purpose of Device Identification and Associated Problems	10
3.5.	Scenario Mapping Table	12
3.6.	Requirements Formulation	13
4.	IANA Considerations	14
5.	Security Considerations	14
6.	Normative References	14
7.	Informative References	14
	Authors' Addresses	15

[1.](#) Introduction

It has become easier for attackers to track the activity of a personal device, particularly when traffic is sent over a wireless link. 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 consent.

To reduce the risks of correlation between a device activity and its owner, multiple vendors have started to implement Randomized and Changing MAC addresses (RCM). With this scheme, an end-device implements a different RCM over time when exchanging traffic over a wireless network. By randomizing the MAC address, the association between a given traffic flow and a single device is made more difficult, assuming no other visible unique identifiers are in use.

However, such address change may affect the user experience and the efficiency of legitimate network operations. For a long time, the unicity of the association between a device and a MAC address was assumed, despite the emergence of tools to flush out the MAC address

to bypass some network policies. When this association is broken, elements of network communication may also break. For example, sessions established between the end-device and network services may be lost and packets in translation may suddenly be without clear source or destination. As multiple clients implements fast-paced RCM rotations, network services may be over-solicited by a small number of stations that appear as many clients.

At the same time, some network services rely on the client station providing an identifier, which can be the MAC address or another value. If the client implements MAC rotation but continues sending the same static identifier, then the association between a stable identifier identifier and the station continues despite the RCM scheme. There may be environments where such continued association is desirable, but others where the user privacy has more value than any continuity of network service state.

There is a need to enumerate services that may be affected by RCM, and evaluate possible solutions to maintain both the quality of user experience and network efficiency while RCM happens and user privacy is reinforced. This document presents such assessment and recommendations.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14 RFC 2119](#) [[RFC2119](#)] [RFC 8174](#) [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

2. MAC Address as an Identity: User vs. Device

Any device member of an IEEE 802 network [[IEEE.802-1D.1993](#)] includes several operating layers. Among them, the Media Access Control (MAC) layer defines rules to control how the 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, either as the target destination of a message, or as the source of a message (and thus the target destination of the answer). Initially intended as a 48-bit (6 octets) value, later versions of the Standard [[IEEE.802.15.4P 2014](#)] allowed this address to take an extended format of 64 bits (8 octets), thus enabling a larger number of MAC addresses to coexist as the 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 either as an individual (bit set to 0) or a group address (bit set to 1). The second bit, called the Universally or Locally Administered (U/L) Address Bit, indicates whether the address has been assigned by a local or universal administrator. Universally administered addresses have this bit set to 0. If this bit is set to 1, the entire address (i.e., 48 bits) has been locally administered [[IEEE.802-1Y.1990](#)] [Section 5.2.1](#).

The intent of this provision is important for the present document. The 802 Standard recognized that some devices may never travel and thus, always attaching to the same network, would not need a globally unique MAC address. To accommodate for this relaxed requirement, the second bit of the MAC address first octet the MAC address format was designed to express whether the address was intended to be globally unique, or if significance was only local. The address allocation method was not defined in the Standard in this later case, but the mechanism was defined in the same clause that defined that an address should be unique so as to avoid collision.

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. The same clause 5.2.1 reminds network designers and operators that all potential members of a network need to have a unique identifier (if they are going to coexist in the network). The advantage of a universal 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.

With the rapid development of wireless technologies and mobile devices, this scenario became very common. With a vast majority of 802 networks implementing radio technologies 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 the 802 Standard generally referred to as 'nodes in a network'. Their definition is found in the 802E Recommended Practice (clause 6.2). One type is a shared service device, which functions are used by a number of people large enough 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, 802.11 (WLAN) access points in a crowded airport, task-specific (e.g. barcode scanners) devices, etc. Another type is a personal device, which is a machine, a node, primarily used by a

single person or small group of people, and so that any identification of the device or its traffic can also be associated to the identification of the primary user or their traffic. Quite naturally, the unique identification of the device is trivial if the device expresses a universally unique MAC address. Then, the detection of elements directly or indirectly identifying the user of the device (Personally Identifiable Information, or PII) is sufficient to tie the universal MAC address to a user. Then, any detection of traffic that can be associated to the device becomes also associated with the known user of that device (Personally Correlated Information, or PCI).

This possible identification or association presents a serious privacy issue, especially with wireless technologies. For most of them, and in particular for 802.11, the source and destination MAC addresses are not encrypted even in networks that implement encryption (so that each machine can easily detect if it is the intended target of the message before attempting to decrypt its content, and also identify the transmitter, so as to use the right key when multiple unicast keys are in effect).

This unique identification of the user associated to a node was clearly not the intent of the 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 temporal association between one MAC address and some PII to be maintained fruitfully. However, other network devices on the same LAN implementing a MAC layer also expect the unicity of the MAC address. 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. Additionally, multiple layers implemented at upper OSI layers have been designed with the assumption that each node on the LAN, using these services, would have a unique MAC address. This assumption sometimes adds to the PII confusion, for example in the case of Authentication, Association and Accounting (AAA) services 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. DHCP), but still expect each device to use a single MAC address. Changing the MAC address may disrupt these services.

3. The Actors: Network Functional Entities and Human Entities

The risk of service disruption is thus weighted against the privacy benefits. However, the plurality of actors involved in the exchanges tends to blurry the boundaries of what privacy should be protected against. It might therefore be useful to list the actors to the network exchanges. Some actors are functional entities, some others are humans (or related) entities.

3.1. Network Functional Entities

Wireless access network infrastructure devices (e.g. WLAN access points or controllers): these devices participate in 802 LAN operations. As such, they need to uniquely identify machines as a source or destination so as to successfully continue exchanging frames. Part of the identification includes recording, and adapting to, devices communication capabilities (e.g. support for specific protocols). 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, 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 OSI model Layer 2 frames are redirected to the new device access point.

Other network devices operating at the MAC layer: many wireless network access devices (e.g., 802.11 access points) are conceived as Layer 2 devices, and as such they bridge a frame from one medium (e.g., 802.11 or Wi-Fi) to another (e.g., 802.3 or Ethernet). This means that a wireless device MAC address often exists on the wire beyond the wireless access device. Devices connected to this wire also implement 802 technologies, and as such operate on the expectation that each device is associated to a unique MAC address for the duration of continuous exchanges. For example, switches and bridges associate MAC addresses to individual ports (so as to know which port to send a frame intended for a particular MAC address). Similarly, authentication, authorization and accounting (AAA) services can validate the identity of a device and use the device MAC address as a 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. 802.1X-enabled devices may also selectively block the data portion of a port until a connecting device is authenticated. These services then use the MAC address as a first pointer to the device identity to allow or block data traffic. This list is not exhaustive. Multiple services are defined for 802.3 networks, and multiple services defined by the IEEE 802.1 working group are also applicable to 802.3 networks. Wireless access points may also

connect to other mediums than 802.3, which also implements mechanism under the umbrella of the general 802 Standard, and therefore expect the unique association of a MAC address to a device.

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 802 principle of unique MAC-to-device mapping. For example, DHCPv4 services commonly provide a single IP address per MAC address (they do not assign more than one IPv4 address per MAC address, and assign a new IPv4 address to each new requesting MAC address). ARP and reverse-ARP services commonly expect that, once an IP-to-MAC mapping has been established, this mapping is valid and unlikely to change for the cache lifetime. DHCPv6 services commonly do not assign the same IPv6 address to two different requesting MAC addresses. Hybrid services, such as EoIP, also assume stability of the device-to-MAC-and-IP mapping for the duration of a given session.

3.2. Human-related Entities

Over the air (OTA) observers: as the transmitting or receiving MAC address is usually not encrypted in wireless 802-technologies exchanges, and as any protocol-compatible device in range of the signal can read the frame header, OTA observers are able to read individual transmissions MAC addresses. Some wireless technologies also support techniques to establish distances or positions, allowing the observer, in some cases, to uniquely associate the MAC address to a physical device and its associated location. It can happen that an OTA observer has a legitimate reason to monitor a particular device, for example for IT support operations. However, it is difficult to control if another actor also monitors the same station with the goal of obtaining PII or PCI.

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

Network access providers: wireless access networks are often considered beyond the first 2 layers of the OSI model. For example, several regulatory or legislative bodies can group all OSI layers into their functional effect of allowing network communication between machines. In this context, entities operating access

networks can see their liability associated to the activity of devices communicating through the networks that these entities operate. In other contexts, operators assign network resources based on contractual conditions (e.g., fee, bandwidth fair share). In these scenarios, these operators may attempt to identify the devices and the users of their networks. They can use the MAC address to represent an identified device.

Over the wire 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 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. 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.

Over the wired external (OTWe) observers: beyond the broadcast domain, frames 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 personal 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 onto the next segment (e.g. pre-[\[RFC4941\]](#) and pre-[\[RFC7217\]](#) IPv6 addresses built from the MAC address). Therefore, unless this last condition exists, OTWe observers are not able to see the device MAC address.

3.3. The Trust and the Environments

The surface of PII exposures that can drive MAC address randomization depends on the environment where the device operates, on the presence and nature of other devices in the environment, and on the type of network the device is communicating through. Therefore, a device can express an identity (such as a MAC address) that can be stable over time if trust with the environment is established, or that can be temporal if an identity is required for a service in an environment where trust has not been established. Trust is not a binary currency. Thus it is useful to distinguish what trust a personal device may establish with the different entities at play in a L2 domain:

1. Full trust: there are environments where a personal device establishes a trust relationship and can share a stable device identity with the access network devices (e.g., access point and WLC), the services beyond the access point in the L2 broadcast domain (e.g. DHCP, AAA). The personal device (or its user) has confidence that its identity is not shared beyond the L2 broadcast domain boundary.
2. Selective trust: in other environments, the device may not be willing to share a stable identity with some elements of the Layer 2 broadcast domain, but may be willing to share a stable identity with other elements. For example, a device may want to change the MAC address it uses to communicate with the access point while maintaining the same IP address across the MAC address rotation (thus expressing a stable identity to the DHCP server). That stable identity may or may not be the same for different services.
3. Zero trust: in other environments, the device may not be willing to share any stable identity with any entity reachable through the Layer 2 broadcast domain, and may express a temporal identity to each of them. That temporal identity may or not be the same for different services.

This trust relationship naturally depends on the relationship between the user of the personal device and the operator of the service. Thus, it is useful to observe the typical trust structure of common environments:

- A. Residential settings under the control of the user: this is typical of a home network with Wi-Fi in the LAN and Internet connection. In this environment, the MAC address activity may be detectable beyond the home walls. However, if traffic is encrypted (e.g. WPA3), some protection for OTA eavesdropping can be assumed. The wire segment within the broadcast domain is under the control of the user, and is therefore usually not at risk of hosting an eavesdropper. Full trust is typically established at this level. The device trusts the access point and all L2 domain entities beyond the access point. Traffic over the Internet does not expose the MAC address if it is not copied to another field before routing happens.
- 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 that of a home. A number of devices larger than in a standard home may be present, and the operator may be requested to provide IT support to the

residents. Therefore, the operator may need to identify a device activity in real time, but may also need to analyze logs so as to understand a past reported issue. For both activities, a device identification associated to the session is needed. Full trust is often established in this environment, at the scale of a series of a few sessions.

- C. Public guest networks: public hotspots, such as in shopping malls, hotels, stores, trains stations and airports are typical of this environment. The guest network operator may be legally mandated to identify devices or users or may have the option to leave all devices and users untracked. In this environment, trust is commonly not established with any element of the L2 broadcast domain (Zero trust model by default).
- D. Enterprises (with BYOD): users may be provided corporate devices or may bring their own devices. The devices are not directly under the control of a corporate IT team. Trust may be established as the device joins the network. Some enterprise models will mandate full trust, others, considering the BYOD nature of the device, will allow selective trust.
- E. Managed enterprises: in this environment, users are typically provided with corporate devices, and all connected devices are managed, for example through a Mobile Device Management (MDM) profile installed on the device. Full trust is created as the MDM profile is installed.

3.4. The Purpose of Device Identification and Associated Problems

Many network functional devices offering a service to a personal device use the device MAC address to maintain service continuity.

Wireless access points and controllers use the MAC address to validate the device connection context, including protocol capabilities, confirmation that authentication was completed, QoS or security profiles, encryption key material. Some advanced access points and controllers also include upper layer functions which 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 these parameters. 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 decrypting the device traffic, and fail selecting the right key to send encrypted traffic 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.) The 802.11i Standard recognizes that a device may leave the network and come back after a short time window. As such, the standard suggests that the infrastructure should keep the context for a device up to one hour after the device was last seen. MAC address rotation in this context can cause resource exhaustion on the wireless infrastructure and the flush of contexts, including for devices that are simply in temporal sleep mode.

Other devices in the Layer 2 broadcast domain also use the MAC address to know whether and where to forward frames. MAC rotation can cause these devices to exhaust their resources, holding in memory traffic for a device which port location can no longer be found. As these infrastructure devices also implement a cache (to remember the port position of each known device), too frequent MAC rotation can cause resources exhaustion and the flush of older MAC addresses, including for devices that did not rotate their MAC. For the RCM device, these effects translate into session discontinuity and return traffic losses.

In wireless contexts, 802.1X authenticators rely on the device and user identity validation provided by a AAA server to open their port to data transmission. The MAC address is used to verify that the device is in the authorized list, and 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. Therefore, MAC rotation can interrupt the device traffic, and cause a strain on the AAA server.

DHCP servers, without a unique identification of the device, lose track of which IP address is validly assigned. Unless the RCM device releases the IP address before the rotation occurs, DHCP 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 the rotation occurs, the DHCP server typically holds the released IP address for a certain duration, in case the leaving MAC would return. As the DHCP server cannot know if the release is due to a temporal disconnection or a MAC rotation, the risk of scope address exhaustion exists even in cases where the IP address is released.

Routers keep track of which MAC address is on which interface. MAC rotation can cause MAC address cache exhaustion, but also the need for frequent ARP and inverse ARP exchanges.

In residential settings (environments type A), policies can be in place to control the traffic of some devices (e.g. parental control, block-list devices). These policies are often based on the device MAC address. Rotation of the MAC address removes the possibility for such control.

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

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

3.5. Scenario Mapping Table

[Section 3.4](#) discusses different environments, different settings, and the expectations of users and network operators. Table 1 summarizes the expected degree of trust, network admin responsibility, complexity of supported network services and network support expectation from the user

Network Location	Trust Degree	Network Admin	Network Services	Network Support Expectation
Home	High	User	Medium	Low
Campus (BYOD)	Medium	IT	Complex	Medium
Enterprise (MDM)	High	IT	Complex	High
Hospitality	Low	IT	Simple	Medium
Public WiFi	Low	ISP	Simple	Low

Table 1: Scenario Mapping Table

For example: a Home network is considered to be trusted and safe. Users expect a simple procedure to connect to their home network. All devices in the home network trust each other. Privacy within the Layer 2 domain is not a major concern. The Home network can also include many IoT devices, which need to be simple to onboard and manage. The home user commonly expects the network operator to protect the home network from external threats (attacks from the Internet). The home user also commonly expects simple policy features (e.g., Parental Control). Most home users do not expect to need networking skills to manage their home network.

On the other end of the spectrum, Public Wi-Fi is often considered to be untrusted. Privacy is the number one concern for the user. Most users connect to Public Wi-Fi only require simple Internet connectivity service, and expect only limited to no technical support.

3.6. Requirements Formulation

The section describes the requirements for Randomized MAC-Address Changes:

- REQ1 The network must not make any assumption about client MAC address persistence. MAC address change must happen while allowing for service continuity. If a service is interrupted during the RCM process, there must be a formal mechanism for the client and the network to exchange about the interruption.
- REQ2 During duration of the services, the device should not change the identity. Any change of identity may result re-authentication and interrupt of the current network services.
- REQ3 Survey the current standards that use MAC address as a device identifier in the protocol. Make recommendation to the working groups to remove the dependency.
- REQ4 Work as liaison with external standard bodies such as IEEE, BBF and WBA to align with use cases and requirements.
- REQ5 Identify a secure mechanism to authenticate and exchange network identity to the device.
- REQ6 Identify a secure mechanism to inform the device about the type of network the device is connecting to (e.g. public Wi-Fi, enterprise, home), allowing the user to select the device identity (or identities) accordingly.

REQ7 Identify a secure mechanism for the network to request device identity. Upon successful authentication, the network may provide the device a temporary network-based marker to use the network services.

REQ8 Identify a secure mechanism for the device to notify the network prior to update the MAC address.

4. IANA Considerations

This memo includes no request to IANA.

5. Security Considerations

Privacy considerations are discussed throughout this document.

6. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", [BCP 72](#), [RFC 3552](#), DOI 10.17487/RFC3552, July 2003, <<https://www.rfc-editor.org/info/rfc3552>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [RFC 5226](#), DOI 10.17487/RFC5226, May 2008, <<https://www.rfc-editor.org/info/rfc5226>>.

7. Informative References

- [IEEE.802-1D.1993]
Institute of Electrical and Electronics Engineers, "Information technology - Telecommunications and information exchange between systems - Local area networks - Media access control (MAC) bridges", IEEE Standard 802.1D, July 1993.
- [IEEE.802-1Y.1990]
Institute of Electrical and Electronics Engineers, "Source Routing Tutorial for End System Operation", IEEE Standard 802.1Y, September 1990.

[IEEE.802.15.4P_2014]

IEEE, "IEEE Standard for local and metropolitan area networks - Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs) - Amendment 7: Physical Layer for Rail Communications and Control (RCC)", IEEE 802.15.4p-2014, DOI 10.1109/ieeestd.2014.6809836, 2 May 2014, <<http://ieeexplore.ieee.org/servlet/opac?punumber=6809834>>.

[RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", [RFC 4941](#), DOI 10.17487/RFC4941, September 2007, <<https://www.rfc-editor.org/info/rfc4941>>.

[RFC5176] Chiba, M., Domety, G., Eklund, M., Mitton, D., and B. Aboba, "Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)", [RFC 5176](#), DOI 10.17487/RFC5176, January 2008, <<https://www.rfc-editor.org/info/rfc5176>>.

[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>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

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: yiulee@comcast.com

