                   Anonymity Profiles for DHCP Clients

Abstract

   Some DHCP options carry unique identifiers. These identifiers can
   enable device tracking even if the device administrator takes care of
   randomizing other potential identifications like link-layer addresses
   or IPv6 addresses. The anonymity profiles are designed for clients
   that wish to remain anonymous to the visited network. The profiles
   provide guidelines on the composition of DHCP or DHCPv6 messages,
   designed to minimize disclosure of identifying information.

Status of This Memo

   This is an Internet Standards Track document.

   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). Further information on
   Internet Standards is available in Section 2 of RFC 5741.

   Information about the current status of this document, any errata,
   and how to provide feedback on it may be obtained at
   http://www.rfc-editor.org/info/rfc7844.

Copyright Notice

Table of Contents

1. Introduction

   There have been reports of systems that would monitor the wireless
   connections of passengers at Canadian airports [CNBC]. We can assume
   that these are either fragments or trial runs of a wider system that
   would attempt to monitor Internet users as they roam through wireless
   access points and other temporary network attachments. We can also
   assume that privacy-conscious users will attempt to evade this
   monitoring -- for example, by ensuring that low-level identifiers
   such as link-layer addresses are "randomized", so that the devices
   do not broadcast the same unique identifier in every location that
   they visit.

   Of course, link-layer MAC (Media Access Control) addresses are not
   the only way to identify a device. As soon as it connects to a
   remote network, the device may use DHCP and DHCPv6 to obtain network
   parameters. The analysis of DHCP and DHCPv6 options shows that
   parameters of these protocols can reveal identifiers of the device,
   negating the benefits of link-layer address randomization. This is
   documented in detail in [RFC7819] and [RFC7824]. The natural
   reaction is to restrict the number and values of such parameters in
   order to minimize disclosure.

   In the absence of a common standard, different system developers are
   likely to implement this minimization of disclosure in different
   ways. Monitoring entities could then use the differences to identify
   the software version running on the device. The proposed anonymity
   profiles provide a common standard that minimizes information
   disclosure, including the disclosure of implementation identifiers.

1.1. Requirements

   The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
   "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
   document are to be interpreted as described in [RFC2119].

2. Application Domain

   Mobile nodes can be tracked using multiple identifiers, the most
   prominent being link-layer addresses, a.k.a. MAC addresses. For
   example, when devices use Wi-Fi connectivity, they place the MAC
   address in the header of all the packets that they transmit.
   Standard implementations of Wi-Fi use unique 48-bit link-layer
   addresses, assigned to the devices according to procedures defined by
   IEEE 802. Even when the Wi-Fi packets are encrypted, the portion of
   the header containing the addresses will be sent in cleartext.
   Tracking devices can "listen to the airwaves" to find out what
   devices are transmitting near them.

We can easily imagine that the MAC addresses can be correlated with
other data, e.g., cleartext names and cookies, to build a registry
linking MAC addresses to the identity of devices' owners. Once that
correlation is done, tracking the MAC address is sufficient to track
individual people, even when all application data sent from the
devices is encrypted. Link-layer addresses can also be correlated
with IP addresses of devices, negating the potential privacy benefits
of IPv6 "privacy" addresses. Privacy advocates have reasons to be
concerned.

The obvious solution is to "randomize" the MAC address. Before
connecting to a particular network, the device replaces the MAC
address with a randomly drawn 48-bit value. Link-layer address
randomization was successfully tried at the IETF meeting in Honolulu
in November 2014 [IETFMACRandom] and in subsequent meetings
[IETFTrialsAndMore]; it is studied in the IEEE 802 EC Privacy
Recommendation Study Group [IEEE802PRSG]. However, we have to
consider the linkage between link-layer addresses, DHCP identifiers,
and IP addresses.

2.1. MAC Address Randomization Hypotheses

There is not yet an established standard for randomizing link-layer
addresses. Various prototypes have tried different strategies,
such as:

Per connection: Configure a random link-layer address at the time of
   connecting to a network, e.g., to a specific Wi-Fi SSID (Service
   Set Identifier), and keep it for the duration of the connection.

Per network: Same as "per connection", but always use the same
   link-layer address for the same network -- different, of course,
   from the addresses used in other networks.

Time interval: Change the link-layer address at regular time
   intervals.

In practice, there are many reasons to keep the link-layer address
constant for the duration of a link-layer connection, as in the
"per connection" or "per network" variants. In Wi-Fi networks,
changing the link-layer address requires dropping the existing Wi-Fi
connection and then re-establishing it, which implies repeating the
connection process and associated procedures. The IP addresses will
change, which means that all required TCP connections will have to be
re-established. If the network access is provided through a NAT,
changing IP addresses also means that the NAT traversal procedures
will have to be restarted. This means a lot of disruption. At the
same time, an observer on the network will easily notice that a

station left, another came in just after that, and the new one
appears to be communicating with the same set of IP addresses as the
old one. This provides for easy correlation.

The anonymity profiles pretty much assume that the link-layer address
randomization follows the "per connection" or "per network"
strategies, or a variant of the "time interval" strategy in which the
interval has about the same duration as the average connection.

2.2. MAC Address Randomization and DHCP

From a privacy point of view, it is clear that the link-layer
address, IP address, and DHCP identifier shall evolve in synchrony.
For example, if the link-layer address changes and the DHCP
identifier stays constant, then it is really easy to correlate old
and new link-layer addresses, either by listening to DHCP traffic or
by observing that the IP address remains constant, since it is tied
to the DHCP identifier. Conversely, if the DHCP identifier changes
but the link-layer address remains constant, the old and new
identifiers and addresses can be correlated by listening to L2
traffic. The procedures documented in the following sections
construct DHCP identifiers from the current link-layer address,
automatically providing for this synchronization.

The proposed anonymity profiles solve this synchronization issue by
deriving most identifiers from the link-layer address and by
generally making sure that DHCP parameter values do not remain
constant after an address change.

2.3. Radio Fingerprinting

MAC address randomization solves the trivial monitoring problem in
which someone just uses a Wi-Fi scanner and records the MAC addresses
seen on the air. DHCP anonymity solves the more elaborate scenario
in which someone monitors link-layer addresses and identities used in
DHCP at the access point or DHCP server. But these are not the only
ways to track a mobile device.

Radio fingerprinting is a process that identifies a radio transmitter
by the unique "fingerprint" of its signal transmission, i.e., the
tiny differences caused by minute imperfections of the radio
transmission hardware. This can be applied to diverse types of
radios, including Wi-Fi as described, for example, in
[WiFiRadioFingerprinting]. No amount of link-layer address
randomization will protect against such techniques. Protections may
exist, but they are outside the scope of the present document.

   On the other hand, we should not renounce randomization just because
   radio fingerprinting exists. The radio fingerprinting techniques are
   harder to deploy than just recording link-layer addresses with a
   scanner. Such techniques can only track devices for which the
   fingerprints are known and thus have a narrower scope of application
   than mass monitoring of addresses and DHCP parameters.

2.4. Operating System Fingerprinting

   When a standard like DHCP allows for multiple options, different
   implementers will make different choices for the options that they
   support or the values they choose for the options. Conversely,
   monitoring the options and values present in DHCP messages reveals
   these differences and allows for "operating system fingerprinting",
   i.e., finding the type and version of software that a particular
   device is running. Finding these versions provides some information
   about the device's identity and thus goes against the goal of
   anonymity.

   The design of the anonymity profiles attempts to minimize the number
   of options and the choice of values, in order to reduce the
   possibilities of operating system fingerprinting.

2.5. No Anonymity Profile Identification

   Reviewers of the anonymity profiles have sometimes suggested adding
   an option to explicitly identify the profiles as "using the anonymity
   option". One suggestion is that the client tell the server about its
   desire to remain anonymous, so that a willing server could cooperate
   and protect the client's privacy. Another possibility would be to
   use a specific privacy-oriented construct, such as, for example, a
   new type of DHCP Unique Identifier (DUID) for a temporary DUID that
   would be changing over time.

   This is not workable in a large number of cases, as it is possible
   that the network operator (or other entities that have access to the
   operator's network) might be actively participating in surveillance
   and anti-privacy, willingly or not. Declaring a preference for
   anonymity is a bit like walking around with a Guy Fawkes mask. (See
   [GuyFawkesMask] for an explanation of this usage.) When anonymity is
   required, it is generally not a good idea to stick out of the crowd.
   Simply revealing the desire for privacy could cause the attacker to
   react by triggering additional surveillance or monitoring mechanisms.
   Therefore, we feel that it is preferable to not disclose one's desire
   for privacy.

   This preference leads to some important implications. In particular,
   we make an effort to make the mitigation techniques difficult to
   distinguish from regular client behaviors, if at all possible.

2.6. Using the Anonymity Profiles

   There are downsides to randomizing link-layer addresses and DHCP
   identifiers. By definition, randomization will break management
   procedures that rely on tracking link-layer addresses. Even if this
   is not too much of a concern, we have to be worried about the
   frequency of link-layer address randomization. Suppose, for example,
   that many devices would get new random link-layer addresses at short
   intervals, maybe every few minutes. This would generate new DHCP
   requests in rapid succession, with a high risk of exhausting DHCPv4
   address pools. Even with IPv6, there would still be a risk of
   increased neighbor discovery traffic and bloating of various address
   tables. Implementers will have to be cautious when programming
   devices to use randomized MAC addresses. They will have to carefully
   choose the frequency with which such addresses will be renewed.

   This document only provides guidelines for using DHCP when clients
   care about privacy. We assume that the request for anonymity is
   materialized by the assignment of a randomized link-layer address to
   the network interface. Once that decision is made, the following
   guidelines will avoid leakage of identity in DHCP parameters or in
   assigned addresses.

   There may be rare situations where the clients want to remain
   anonymous to attackers but not to the DHCP server. These clients
   should still use link-layer address randomization to hide from
   observers, as well as some form of encrypted communication to the
   DHCP server. This scenario is out of scope for this document.

   To preserve anonymity, the clients need to not use stable values for
   the client identifiers. This is clearly a trade-off, because a
   stable client identifier guarantees that the client will receive
   consistent parameters over time. An example is given in [RFC7618],
   where the client identifier is used to guarantee that the same client
   will always get the same combination of IP address and port range.
   Static clients benefit most from stable parameters and often can
   already be identified by physical-connection-layer parameters. These
   static clients will normally not use the anonymity profiles. Mobile
   clients, in contrast, have the option of using the anonymity profiles
   in conjunction with [RFC7618] if they are more concerned with privacy
   protection than with stable parameters.

2.7. What about privacy for DHCP servers?

   This document only provides recommendations for DHCP clients. The
   main targets are DHCP clients used in mobile devices. Such devices
   are tempting targets for various monitoring systems, so there is an
   urgent need to provide them with a simple anonymity solution. We can
   argue that some mobile devices embed DHCP servers and that providing
   solutions for such devices is also quite important. Two plausible
   examples would be a DHCP server for a car network and a DHCP server
   for a mobile hot spot. However, mobile servers get a lot of privacy
   protection through the use of access control and link-layer
   encryption. Servers may disclose information to clients through
   DHCP, but they normally only do that to clients that have passed the
   link-layer access control and have been authorized to use the network
   services. This arguably makes solving the server problem less urgent
   than solving the client problem.

   Server privacy issues are presented in [RFC7819] and [RFC7824].
   Mitigation of these issues is left for further study.

3. Anonymity Profile for DHCPv4

   Clients using the DHCPv4 anonymity profile limit the disclosure of
   information by controlling the header parameters and by limiting the
   number and values of options. The number of options depends on the
   specific DHCP message:

   DHCPDISCOVER: The anonymized DHCPDISCOVER messages MUST contain the
      Message Type option, MAY contain the Client Identifier option, and
      MAY contain the Parameter Request List option. It SHOULD NOT
      contain any other option.

   DHCPREQUEST: The anonymized DHCPREQUEST messages MUST contain the
      Message Type option, MAY contain the Client Identifier option, and
      MAY contain the Parameter Request List option. If the message is
      in response to a DHCPOFFER, it MUST contain the corresponding
      Server Identifier option and the Requested IP address option. If
      the message is not in response to a DHCPOFFER, it MAY contain a
      Requested IP address option as explained in Section 3.3. It
      SHOULD NOT contain any other option.

   DHCPDECLINE: The anonymized DHCPDECLINE messages MUST contain the
      Message Type option, the Server Identifier option, and the
      Requested IP address option; and MAY contain the Client Identifier
      option.

   DHCPRELEASE: The anonymized DHCPRELEASE messages MUST contain the
      Message Type option and the Server Identifier option, and MAY
      contain the Client Identifier option.

   DHCPINFORM: The anonymized DHCPINFORM messages MUST contain the
      Message Type option, MAY contain the Client Identifier option, and
      MAY contain the Parameter Request List option. It SHOULD NOT
      contain any other option.

   Header fields and option values SHOULD be set in accordance with the
   DHCP specification, but some header fields and option values SHOULD
   be constructed per the following guidelines.

   The inclusion of the Host Name and Fully Qualified Domain Name (FQDN)
   options in DHCPDISCOVER, DHCPREQUEST, or DHCPINFORM messages is
   discussed in Sections 3.7 and 3.8.

3.1. Avoiding Fingerprinting

   There are many choices for implementing DHCPv4 messages. Clients can
   choose to transmit a specific set of options, pick a particular
   encoding for these options, and transmit options in different orders.
   These choices can be used to fingerprint the client.

   The following sections provide guidance on the encoding of options
   and fields within the packets. However, this guidance alone may not
   be sufficient to prevent fingerprinting from revealing the device
   type, the vendor name, or the OS type and specific version.
   Fingerprinting may also reveal whether the client is using the
   anonymity profile.

   The client intending to protect its privacy SHOULD limit the subset
   of options sent in messages to the subset listed in the remaining
   subsections.

   The client intending to protect its privacy SHOULD randomize the
   ordering of options before sending any DHCPv4 message. If this
   random ordering cannot be implemented, the client MAY order the
   options by option code number (lowest to highest).

3.2. Client IP Address Field

   Four octets in the header of the DHCP messages carry the "Client IP
   address" (ciaddr) as defined in [RFC2131]. In DHCP, this field is
   used by the clients to indicate the address that they used
   previously, so that as much as possible the server can allocate the
   same address to them.

There are very few privacy implications related to sending this
address in the DHCP messages, except in the case of connecting to a
different network than the last network connected to previously. If
the DHCP client somehow repeated the address used in a previous
network attachment, monitoring services might use the information to
tie the two network locations. DHCP clients SHOULD ensure that the
field is cleared when they know that the network attachment has
changed, particularly if the link-layer address is reset by a
device's administrator.

The clients using the anonymity profile MUST NOT include in the
message a Client IP address that has been obtained with a different
link-layer address.

3.3. Requested IP Address Option

The Requested IP address option is defined in [RFC2132] with code 50.
It allows the client to request that a particular IP address be
assigned. This option is mandatory in some protocol messages per
[RFC2131] -- for example, when a client selects an address offered by
a server. However, this option is not mandatory in the DHCPDISCOVER
message. It is simply a convenience -- an attempt to regain the same
IP address that was used in a previous connection. Doing so entails
the risk of disclosing an IP address used by the client at a previous
location or with a different link-layer address. This risk exists
for all forms of IP addresses, public or private, as some private
addresses may be used in a wide scope, e.g., when an Internet Service
Provider is using NAT.

When using the anonymity profile, clients SHOULD NOT use the
Requested IP address option in DHCPDISCOVER messages. They MUST use
the option when mandated by DHCP -- for example, in DHCPREQUEST
messages.

There are scenarios in which a client connecting to a network
remembers a previously allocated address, i.e., when it is in the
INIT-REBOOT state. In that state, any client that is concerned with
privacy SHOULD perform a complete four-way handshake, starting with a
DHCPDISCOVER, to obtain a new address lease. If the client can
ascertain that this is exactly the same network to which it was
previously connected, and if the link-layer address did not change,
the client MAY issue a DHCPREQUEST to try to reclaim the current
address.

3.4. Client Hardware Address Field

   Sixteen octets in the header of the DHCP messages carry the "Client
   hardware address" (chaddr) as defined in [RFC2131]. The presence of
   this address is necessary for the proper operation of the DHCP
   service.

   Hardware addresses, called "link-layer addresses" in many RFCs, can
   be used to uniquely identify a device, especially if they follow the
   IEEE 802 recommendations. If the hardware address is reset to a new
   randomized value, the DHCP client SHOULD use the new randomized value
   in the DHCP messages.

3.5. Client Identifier Option

   The Client Identifier option is defined in [RFC2132] with
   option code 61. It is discussed in detail in [RFC4361]. The purpose
   of the Client Identifier option is to identify the client in a manner
   independent of the link-layer address. This is particularly useful
   if the DHCP server is expected to assign the same address to the
   client after a network attachment is swapped and the link-layer
   address changes. It is also useful when the same node issues
   requests through several interfaces and expects the DHCP server to
   provide consistent configuration data over multiple interfaces.

   The considerations for hardware independence and strong client
   identity have an adverse effect on the privacy of mobile clients,
   because the hardware-independent unique identifier obviously enables
   very efficient tracking of the clients' movements. One option would
   be to not transmit this option at all, but this may affect
   interoperability and will definitely mark the client as requesting
   anonymity, exposing it to the risks mentioned in Section 2.5.

   The recommendations in [RFC4361] are very strong, stating, for
   example, that "DHCPv4 clients MUST NOT use client identifiers based
   solely on layer two addresses that are hard-wired to the layer two
   device (e.g., the Ethernet MAC address)." These strong
   recommendations are in fact a trade-off between ease of management
   and privacy, and the trade-off should depend on the circumstances.

   In contradiction to [RFC4361], when using the anonymity profile, DHCP
   clients MUST use client identifiers based solely on the link-layer
   address that will be used in the underlying connection. This will
   ensure that the DHCP client identifier does not leak any information
   that is not already available to entities monitoring the network
   connection. It will also ensure that a strategy of randomizing the
   link-layer address will not be nullified by the Client Identifier
   option.

   There are usages of DHCP where the underlying connection is a
   point-to-point link, in which case there is no link-layer address
   available to construct a non-revealing identifier. If anonymity is
   desired in such networks, the client SHOULD pick a random identifier
   that is highly likely to be unique to the current link, using, for
   example, a combination of a local secret and an identifier of the
   connection. The algorithm for combining secrets and identifiers, as
   described in Section 5 of [RFC7217], solves a similar problem. The
   criteria for the generation of random numbers are stated
   in [RFC4086].

3.6. Parameter Request List Option

   The Parameter Request List (PRL) option is defined in [RFC2132] with
   option code 55. It lists the parameters requested from the server by
   the client. Different implementations request different parameters.
   [RFC2132] specifies that "the client MAY list the options in order of
   preference." In practice, this means that different client
   implementations will request different parameters, in different
   orders.

   The choice of option numbers and the specific ordering of option
   numbers in the PRL can be used to fingerprint the client. This may
   not reveal the identity of a client but may provide additional
   information such as the device type, the vendor name, or the OS type
   and specific version.

   The client intending to protect its privacy SHOULD only request a
   minimal number of options in the PRL and SHOULD also randomly shuffle
   the ordering of option codes in the PRL. If this random ordering
   cannot be implemented, the client MAY order the option codes in the
   PRL by option code number (lowest to highest).

3.7. Host Name Option

   The Host Name option is defined in [RFC2132] with option code 12.
   Depending on implementations, the option value can carry either an
   FQDN such as "node1984.example.com" or a simple host name such as
   "node1984". The host name is commonly used by the DHCP server to
   identify the host and also to automatically update the address of the
   host in local name services.

   FQDNs are obviously unique identifiers, but even simple host names
   can provide a significant amount of information on the identity of
   the device. They are typically chosen to be unique in the context
   where the device is most often used. In a context that contains a
   substantial number of devices, e.g., in a large company or a big
   university, the host name will be a pretty good identifier of the

   device, due to the specificity required to ensure uniqueness.
   Monitoring services could use that information in conjunction with
   traffic analysis and quickly derive the identity of the device's
   owner.

   When using the anonymity profile, DHCP clients SHOULD NOT send the
   Host Name option. If they choose to send the option, DHCP clients
   MUST always send a non-qualified host name instead of an FQDN and
   MUST obfuscate the host name value.

   There are many ways to obfuscate a host name. The construction rules
   SHOULD guarantee that a different host name is generated each time
   the link-layer address changes and that the obfuscated host name will
   not reveal the underlying link-layer address. The construction
   SHOULD generate names that are unique enough to minimize collisions
   in the local link. Clients MAY use the following algorithm: compute
   a secure hash of a local secret and of the link-layer address that
   will be used in the underlying connection, and then use the
   hexadecimal representation of the first 6 octets of the hash as the
   obfuscated host name.

   The algorithm described in the previous paragraph generates an easily
   recognizable pattern. There is a potential downside to having such a
   specific name pattern for hosts that require anonymity (the "sticking
   out of the crowd" principle), as explained in Section 2.5. For this
   reason, the above algorithm is just a suggestion.

3.8. Client FQDN Option

   The Client FQDN option is defined in [RFC4702] with option code 81.
   This option allows the DHCP clients to advertise to the DHCP server
   their FQDN, such as "mobile.example.com". This would allow the DHCP
   server to update in the DNS the PTR record for the IP address
   allocated to the client. Depending on circumstances, either the DHCP
   client or the DHCP server could update in the DNS the A record for
   the FQDN of the client.

   Obviously, this option uniquely identifies the client, exposing it to
   the DHCP server or to anyone listening to DHCP traffic. In fact, if
   the DNS record is updated, the location of the client becomes visible
   to anyone with DNS lookup capabilities.

   When using the anonymity profile, DHCP clients SHOULD NOT include the
   Client FQDN option in their DHCP requests. Alternatively, they MAY
   include a special-purpose FQDN using the same host name as in the
   Host Name option, with a suffix matching the connection-specific DNS
   suffix being advertised by that DHCP server. Having a name in the

DNS allows working with legacy systems that require one to be there,
e.g., by verifying that a forward and reverse lookup succeeds with
the same result.

3.9. UUID/GUID-Based Client Machine Identifier Option

The UUID/GUID-based (where "UUID" means "Universally Unique
Identifier" and "GUID" means "Globally Unique Identifier")
Client Machine Identifier option is defined in [RFC4578] with
option code 97. This option is part of a set of options for the
Intel Preboot eXecution Environment (PXE). The purpose of the PXE
system is to perform management functions on a device before its main
OS is operational. The Client Machine Identifier carries a 16-octet
GUID that uniquely identifies the device.

The PXE system is clearly designed for devices operating in a
controlled environment. The main usage of the PXE system is to
install a new version of the operating system through a high-speed
Ethernet connection. The process is typically controlled from the
user interface during the boot process. Common sense seems to
dictate that getting a new operating system from an unauthenticated
server at an untrusted location is a really bad idea and that even if
the option was available users would not activate it. In any case,
the option is only used in the "pre-boot" environment, and there is
no reason to use it once the system is up and running. Nodes
visiting untrusted networks MUST NOT send or use the PXE options.

3.10. User and Vendor Class DHCP Options

Vendor-identifying options are defined in [RFC2132] and [RFC3925].
When using the anonymity profile, DHCPv4 clients SHOULD NOT use the
Vendor-Specific Information option (code 43), the Vendor Class
Identifier option (code 60), the V-I Vendor Class option (code 124),
or the V-I Vendor-Specific Information option (code 125), as these
options potentially reveal identifying information.

4. Anonymity Profile for DHCPv6

DHCPv6 is typically used by clients in one of two scenarios: stateful
or stateless configuration. In the stateful scenario, clients use a
combination of Solicit, Request, Confirm, Renew, Rebind, Release, and
Decline messages to obtain addresses and manage these addresses.

In the stateless scenario, clients configure addresses using a
combination of client-managed identifiers and router-advertised
prefixes, without involving the DHCPv6 services. Different ways of
constructing these prefixes have different implications on privacy,
which are discussed in [DEFAULT-IIDs] and [RFC7721]. In the
stateless scenario, clients use DHCPv6 to obtain network
configuration parameters, through the Information-request message.

The choice between the stateful and stateless scenarios depends on
flag and prefix options published by the Router Advertisement
messages of local routers, as specified in [RFC4861]. When these
options enable stateless address configuration, hosts using the
anonymity profile SHOULD use stateless address configuration instead
of stateful address configuration, because stateless configuration
requires fewer information disclosures than stateful configuration.

When using the anonymity profile, DHCPv6 clients carefully select
DHCPv6 options used in the various messages that they send. The list
of options that are mandatory or optional for each message is
specified in [RFC3315]. Some of these options have specific
implications on anonymity. The following sections provide guidance
on the choice of option values when using the anonymity profile.

4.1. Avoiding Fingerprinting

There are many choices for implementing DHCPv6 messages. As
explained in Section 3.1, these choices can be used to fingerprint
the client.

The following sections provide guidance on the encoding of options.
However, this guidance alone may not be sufficient to prevent
fingerprinting from revealing the device type, the vendor name, or
the OS type and specific version. Fingerprinting may also reveal
whether the client is using the anonymity profile.

The client intending to protect its privacy SHOULD limit the subset
of options sent in messages to the subset listed in the following
sections.

The client intending to protect its privacy SHOULD randomize the
ordering of options before sending any DHCPv6 message. If this
random ordering cannot be implemented, the client MAY order the
options by option code number (lowest to highest).

4.2. Do not send Confirm messages, unless really sure about
      the location

   [RFC3315] requires clients to send a Confirm message when they attach
   to a new link to verify whether the addressing and configuration
   information they previously received is still valid. This
   requirement was relaxed in [DHCPv6bis]. When these clients send
   Confirm messages, they include any Identity Associations (IAs)
   assigned to the interface that may have moved to a new link, along
   with the addresses associated with those IAs. By examining the
   addresses in the Confirm message, an attacker can trivially identify
   the previous point(s) of attachment.

   Clients interested in protecting their privacy SHOULD NOT send
   Confirm messages and instead SHOULD directly try to acquire addresses
   on the new link. However, not sending Confirm messages can result in
   connectivity hiatus in some scenarios, e.g., roaming between two
   access points in the same wireless network. DHCPv6 clients that can
   verify that the previous link and the current link are part of the
   same network MAY send Confirm messages while still protecting their
   privacy. Such link identification should happen before DHCPv6 is
   used, and thus it cannot depend on the DHCPv6 information used in
   [RFC6059]. In practice, the most reliable detection of network
   attachment is through link-layer security, e.g., [IEEE8021X].

4.3. Client Identifier DHCPv6 Option

   The DHCPv6 Client Identifier option is defined in [RFC3315] with
   option code 1. The purpose of the Client Identifier option is to
   identify the client to the server. The content of the option is a
   DHCP Unique Identifier (DUID). One of the primary privacy concerns
   is that a client is disclosing a stable identifier (the DUID) that
   can be used for tracking and profiling. Three DUID formats are
   specified in [RFC3315]: link-layer address plus time (DUID-LLT),
   Vendor-assigned unique ID based on Enterprise Number, and link-layer
   address. A fourth type, DUID-UUID, is defined in [RFC6355].

   When using the anonymity profile in conjunction with randomized
   link-layer addresses, DHCPv6 clients MUST use DUID format number 3 --
   link-layer address. The value of the link-layer address should be
   the value currently assigned to the interface.

   When using the anonymity profile without the benefit of randomized
   link-layer addresses, clients that want to protect their privacy
   SHOULD generate a new randomized DUID-LLT every time they attach to a
   new link or detect a possible link change event. Syntactically, this
   identifier will conform to [RFC3315], but its content is meaningless.
   The exact details are left up to implementers, but there are several

factors that should be taken into consideration. The DUID type
SHOULD be set to 1 (DUID-LLT). Hardware type SHOULD be set
appropriately to the hardware type in question. The link address
embedded in the LLT SHOULD be set to a randomized value. Time SHOULD
be set to a random timestamp from the previous year. Time MAY be set
to current time, but this will reveal the fact that the DUID is newly
generated and thus could provide information for device
fingerprinting. The criteria for generating highly unique random
numbers are listed in [RFC4086].

4.3.1. Anonymous Information-request

   According to [RFC3315], a DHCPv6 client includes its client
   identifier in most of the messages it sends. There is one exception,
   however: the client is allowed to omit its client identifier when
   sending Information-request messages.

   When using stateless DHCPv6, clients wanting to protect their privacy
   SHOULD NOT include client identifiers in their Information-request
   messages. This will prevent the server from specifying client-
   specific options if it is configured to do so, but the need for
   anonymity precludes such options anyway.

4.4. Server Identifier Option

   When using the anonymity profile, DHCPv6 clients SHOULD use the
   Server Identifier option (code 2) as specified in [RFC3315]. Clients
   MUST only include server identifier values that were received with
   the current link-layer address, because the reuse of old values
   discloses information that can be used to identify the client.

4.5. Address Assignment Options

   When using the anonymity profile, DHCPv6 clients might have to use
   Solicit or Request messages to obtain IPv6 addresses through the
   DHCPv6 server. In DHCPv6, the collection of addresses assigned to a
   client is identified by an IA. Clients interested in privacy SHOULD
   request addresses using the IA for the Non-temporary Addresses option
   (IA_NA, code 3) [RFC3315].

   The IA_NA option includes an IAID parameter that identifies a unique
   IA for the interface for which the address is requested. Clients
   interested in protecting their privacy MUST ensure that the IAID does
   not enable client identification. They also need to conform to the
   requirement of [RFC3315] that the IAID for that IA MUST be consistent
   across restarts of the DHCPv6 client. We interpret that as requiring
   that the IAID MUST be constant for the association, as long as the
   link-layer address remains constant.

Clients MAY meet the privacy, uniqueness, and stability requirements
of the IAID by constructing it as the combination of 1 octet encoding
the interface number in the system, and the first 3 octets of the
link-layer address.

The clients MAY use the IA Address option (code 5) [RFC3315] but need
to balance the potential advantage of "address continuity" versus the
potential risk of "previous address disclosure". A potential
solution is to remove all stored addresses when a link-layer address
changes and to only use the IA Address option with addresses that
have been explicitly assigned through the current link-layer address.

4.5.1. Obtain Temporary Addresses

[RFC3315] defines a special container (IA_TA, code 4) for requesting
temporary addresses. This is a good mechanism in principle, but
there are a number of issues associated with it. First, this is not
a widely used feature, so clients depending solely on temporary
addresses may lock themselves out of service. Secondly, [RFC3315]
does not specify any lifetime or lease length for temporary
addresses. Therefore, support for renewing temporary addresses may
vary between client implementations, including no support at all.
Finally, by requesting temporary addresses, a client reveals its
desire for privacy and potentially risks countermeasures as described
in Section 2.5.

Because of these issues, clients interested in their privacy
SHOULD NOT use IA_TA.

The addresses obtained according to Section 4.5 are meant to be
non-temporary, but the anonymity profile uses them as temporary, and
they will be discarded when the link-layer address is changed. They
thus meet most of the use cases of the temporary addresses defined in
[RFC4941]. Clients interested in their privacy should not publish
their IPv6 addresses in the DNS or otherwise associate them with name
services, and thus do not normally need two classes of addresses --
one public, one temporary.

The use of mechanisms to allocate several IPv6 addresses to a client
while preserving privacy is left for further study.

4.5.2. Prefix Delegation

   The use of DHCPv6 address assignment option for Prefix Delegation
   (PD) is defined in [RFC3633]. Because current host OS
   implementations do not typically request prefixes, clients that wish
   to use DHCPv6 PD -- just like clients that wish to use any DHCP or
   DHCPv6 option that is not currently widely used -- should recognize
   that doing so will serve as a form of fingerprinting, unless or until
   the use of DHCPv6 PD by clients becomes more widespread.

   The anonymity properties of DHCPv6 PD, which uses IA_PD IAs, are
   similar to those of DHCPv6 address assignment using IA_NA IAs. The
   IAID could potentially be used to identify the client, and a prefix
   hint sent in the IA_PD Prefix option could be used to track the
   client's previous location. Clients that desire anonymity and never
   request more than one prefix SHOULD set the IAID value to zero, as
   authorized in Section 6 of [RFC3633], and SHOULD NOT document any
   previously assigned prefix in the IA_PD Prefix option.

4.6. Option Request Option

   The Option Request Option (ORO) is defined in [RFC3315] with
   option code 6. It specifies the options that the client is
   requesting from the server. The choice of requested options and the
   order of encoding of these options in the ORO can be used to
   fingerprint the client.

   The client intending to protect its privacy SHOULD only request a
   minimal subset of options and SHOULD randomly shuffle the ordering of
   option codes in the ORO. If this random ordering cannot be
   implemented, the client MAY order the option codes in the ORO by
   option code number (lowest to highest).

4.6.1. Previous Option Values

   According to [RFC3315], the client that includes an ORO in a Solicit
   or Request message MAY additionally include instances of those
   options that are identified in the ORO, with data values as hints to
   the server about parameter values the client would like to have
   returned.

   When using the anonymity profile, clients SHOULD NOT include such
   instances of options, because old values might be used to identify
   the client.

4.7. Authentication Option

   The purpose of the Authentication option (code 11) [RFC3315] is to
   authenticate the identity of clients and servers and the contents of
   DHCPv6 messages. As such, the option can be used to identify the
   client, so it is incompatible with the stated goal of "client
   anonymity". DHCPv6 clients that use the anonymity profile SHOULD NOT
   use the Authentication option. They MAY use it if they recognize
   that they are operating in a trusted environment, e.g., in a
   workplace network.

4.8. User and Vendor Class DHCPv6 Options

   When using the anonymity profile, DHCPv6 clients SHOULD NOT use the
   User Class option (code 15) or the Vendor Class option (code 16)
   [RFC3315], as these options potentially reveal identifying
   information.

4.9. Client FQDN DHCPv6 Option

   The DHCPv6 Client FQDN option is defined in [RFC4704] with
   option code 39. This option allows the DHCPv6 clients to advertise
   to the DHCPv6 server their FQDN, such as "mobile.example.com". When
   using the anonymity profile, DHCPv6 clients SHOULD NOT include the
   Client FQDN option in their DHCPv6 messages, because it identifies
   the client. As explained in Section 3.8, they MAY use a local-only
   FQDN by combining a host name derived from the link-layer address and
   a suffix advertised by the local DHCPv6 server.

5. Operational Considerations

   The anonymity profiles have the effect of hiding the client identity
   from the DHCP server. This is not always desirable. Some DHCP
   servers provide facilities like publishing names and addresses in the
   DNS, or ensuring that returning clients get reassigned the same
   address.

   Clients using an anonymity profile may be consuming more resources.
   For example, when a client changes its link-layer address and
   requests a new IP address, the old IP address is still marked as
   leased by the server.

   Some DHCP servers will only give addresses to pre-registered MAC
   addresses, forcing clients to choose between remaining anonymous and
   obtaining connectivity.

   Implementers SHOULD provide a way for clients to control when the
   anonymity profiles are used and when standard behavior is preferred.

Implementers MAY implement this control by tying the use of the
anonymity profiles to that of link-layer address randomization.

6. Security Considerations

The use of the anonymity profiles does not change the security
considerations of the DHCPv4 or DHCPv6 protocols [RFC2131] [RFC3315].

7. References

7.1. 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,
              <http://www.rfc-editor.org/info/rfc2119>.

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

   [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins,
              C., and M. Carney, "Dynamic Host Configuration Protocol
              for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315,
              July 2003, <http://www.rfc-editor.org/info/rfc3315>.

   [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic
              Host Configuration Protocol (DHCP) version 6", RFC 3633,
              DOI 10.17487/RFC3633, December 2003,
              <http://www.rfc-editor.org/info/rfc3633>.

   [RFC4702] Stapp, M., Volz, B., and Y. Rekhter, "The Dynamic Host
              Configuration Protocol (DHCP) Client Fully Qualified
              Domain Name (FQDN) Option", RFC 4702,
              DOI 10.17487/RFC4702, October 2006,
              <http://www.rfc-editor.org/info/rfc4702>.

   [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,
              <http://www.rfc-editor.org/info/rfc4861>.

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

7.2. Informative References

   [CNBC]   Weston, G., Greenwald, G., and R. Gallagher, "CBC News:
               CSEC used airport Wi-Fi to track Canadian travellers:
               Edward Snowden documents", January 2014,
               <http://www.cbc.ca/news/politics/
               csec-used-airport-wi-fi-to-track-canadian-travellers-
               edward-snowden-documents-1.2517881>.

   [DEFAULT-IIDs]
               Gont, F., Cooper, A., Thaler, D., and W. Liu,
               "Recommendation on Stable IPv6 Interface Identifiers",
               Work in Progress, draft-ietf-6man-default-iids-11,
               April 2016.

   [DHCPv6bis]
               Mrugalski, T., Ed., Siodelski, M., Volz, B., Yourtchenko,
               A., Richardson, M., Jiang, S., and T. Lemon, "Dynamic Host
               Configuration Protocol for IPv6 (DHCPv6) bis", Work in
               Progress, draft-ietf-dhc-rfc3315bis-04, March 2016.

   [GuyFawkesMask]
               Nickelsburg, M., "A brief history of the Guy Fawkes mask",
               July 2013, <http://theweek.com/articles/463151/
               brief-history-guy-fawkes-mask>.

   [IEEE8021X]
               IEEE, "IEEE Standard for Local and metropolitan area
               networks - Port-Based Network Access Control",
               IEEE 802.1X-2010, DOI 10.1109/ieeestd.2010.5409813,
               <http://ieeexplore.ieee.org/servlet/
               opac?punumber=5409757>.

   [IEEE802PRSG]
               IEEE 802 EC PRSG, "IEEE 802 EC Privacy Recommendation
               Study Group", February 2016,
               <http://www.ieee802.org/PrivRecsg/>.

   [IETFMACRandom]
               Zuniga, JC., "MAC Privacy", November 2014,
               <http://www.ietf.org/blog/2014/11/mac-privacy/>.

   [IETFTrialsAndMore]
               Bernardos, CJ., Zuniga, JC., and P. O'Hanlon, "Wi-Fi
               Internet connectivity and privacy: hiding your tracks on
               the wireless Internet", October 2015,
               <http://www.it.uc3m.es/cjbc/papers/
               pdf/2015_bernardos_cscn_privacy.pdf>.

   [RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor
              Extensions", RFC 2132, DOI 10.17487/RFC2132, March 1997,
              <http://www.rfc-editor.org/info/rfc2132>.

   [RFC3925] Littlefield, J., "Vendor-Identifying Vendor Options for
              Dynamic Host Configuration Protocol version 4 (DHCPv4)",
              RFC 3925, DOI 10.17487/RFC3925, October 2004,
              <http://www.rfc-editor.org/info/rfc3925>.

   [RFC4086] Eastlake 3rd, D., Schiller, J., and S. Crocker,
              "Randomness Requirements for Security", BCP 106, RFC 4086,
              DOI 10.17487/RFC4086, June 2005,
              <http://www.rfc-editor.org/info/rfc4086>.

   [RFC4361] Lemon, T. and B. Sommerfeld, "Node-specific Client
              Identifiers for Dynamic Host Configuration Protocol
              Version Four (DHCPv4)", RFC 4361, DOI 10.17487/RFC4361,
              February 2006, <http://www.rfc-editor.org/info/rfc4361>.

   [RFC4578] Johnston, M. and S. Venaas, Ed., "Dynamic Host
              Configuration Protocol (DHCP) Options for the Intel
              Preboot eXecution Environment (PXE)", RFC 4578,
              DOI 10.17487/RFC4578, November 2006,
              <http://www.rfc-editor.org/info/rfc4578>.

   [RFC4704] Volz, B., "The Dynamic Host Configuration Protocol for
              IPv6 (DHCPv6) Client Fully Qualified Domain Name (FQDN)
              Option", RFC 4704, DOI 10.17487/RFC4704, October 2006,
              <http://www.rfc-editor.org/info/rfc4704>.

   [RFC6059] Krishnan, S. and G. Daley, "Simple Procedures for
              Detecting Network Attachment in IPv6", RFC 6059,
              DOI 10.17487/RFC6059, November 2010,
              <http://www.rfc-editor.org/info/rfc6059>.

   [RFC6355] Narten, T. and J. Johnson, "Definition of the UUID-Based
              DHCPv6 Unique Identifier (DUID-UUID)", RFC 6355,
              DOI 10.17487/RFC6355, August 2011,
              <http://www.rfc-editor.org/info/rfc6355>.

   [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,
              <http://www.rfc-editor.org/info/rfc7217>.

   [RFC7618] Cui, Y., Sun, Q., Farrer, I., Lee, Y., Sun, Q., and M.
              Boucadair, "Dynamic Allocation of Shared IPv4 Addresses",
              RFC 7618, DOI 10.17487/RFC7618, August 2015,
              <http://www.rfc-editor.org/info/rfc7618>.

   [RFC7721] Cooper, A., Gont, F., and D. Thaler, "Security and Privacy
              Considerations for IPv6 Address Generation Mechanisms",
              RFC 7721, DOI 10.17487/RFC7721, March 2016,
              <http://www.rfc-editor.org/info/rfc7721>.

   [RFC7819] Jiang, S., Krishnan, S., and T. Mrugalski, "Privacy
              Considerations for DHCP", RFC 7819, DOI 10.17487/RFC7819,
              April 2016, <http://www.rfc-editor.org/info/rfc7819>.

   [RFC7824] Krishnan, S., Mrugalski, T., and S. Jiang, "Privacy
              Considerations for DHCPv6", RFC 7824,
              DOI 10.17487/RFC7824, May 2016,
              <http://www.rfc-editor.org/info/rfc7824>.

   [WiFiRadioFingerprinting]
              Brik, V., Banerjee, S., Gruteser, M., and S. Oh, "Wireless
              Device Identification with Radiometric Signatures",
              DOI 10.1.1.145.8873, September 2008,
              <http://citeseerx.ist.psu.edu/viewdoc/
              summary?doi=10.1.1.145.8873>.

Authors' Addresses

   Christian Huitema
   Microsoft
   Redmond, WA 98052
   United States

   Email: huitema@microsoft.com

   Tomek Mrugalski
   Internet Systems Consortium, Inc.
   950 Charter Street
   Redwood City, CA 94063
   United States

   Email: tomasz.mrugalski@gmail.com

   Suresh Krishnan
   Ericsson
   8400 Decarie Blvd.
   Town of Mount Royal, QC
   Canada

   Phone: +1 514 345 7900 x42871
   Email: suresh.krishnan@ericsson.com