Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 3, 2016

C. Huitema
Microsoft

March 2, 2016

Implications of Randomized Link Layers Addresses for IPv6 Address Assignment
draft-huitema-6man-random-addresses-03.txt

Abstract

Hosts may assign random link-layer addresses to network interfaces in an attempt to increase privacy and reduce trackability. Careless assignment of IPv6 addresses may negate the privacy advantages of random link-layer addresses. We propose simple solutions to ensure that IPv6 addresses do change whenever the link layer addresses change.

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 http://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 September 3, 2016.

Table of Contents

1. Introduction

The IPv6 Maintenance Working Group is reviewing the privacy properties of various IPv6 address generation mechanisms [I-D.ietf-6man-ipv6-address-generation-privacy]. At the same time, this working group has proposed in [RFC7217] a method for the construction of stable IPv6 identifiers. The method defined in [RFC7217] is designed to prevent address scanning or device identification through the use of "opaque" identifiers. It prevents location tracking by making sure that the same device uses different identifiers at different locations. However, a strict implementation of [RFC7217] results in stable identifiers, which remain always the same for a given device and a given location. This is in fact a design goal of [RFC7217].

Privacy conscious users will not agree with this design goal. Suppose for example users who don't want being tracked when they visit an public place at different times. They will configure their device to use different link layer addresses on the different visits, using a form of MAC Address Randomization, as discussed in Section 2.

However, if their devices implement a strict version of [RFC7217], the IPv6 addresses will contain stable identifiers. The stable identifiers will re-enable the tracking that MAC Address Randomization would have prevented.

Some systems also use temporary IPv6 addresses, as defined by [RFC4941]. These randomized addresses are defined by generating a randomized interface identifier at controlled intervals, and then using this identifier in conjunction with prefixes advertised by routers to construct addresses with limited life time. Even with this short life time, the randomized interface identifier could remain constant while the link layer addresses changes with MAC Address Randomization. This would enable tracking between successive network connections, even if the MAC Address changed.

The purpose of this document is to recommend specific guidelines for the use of [RFC7217] and [RFC4941], in order to make it maintain the privacy benefits of MAC Address Randomization. Section 2 presents the address randomization mechanisms. Section 3 presents the guidelines for use of [RFC7217]. Section 4 presents the guidelines for use of [RFC4941]. Section 5 reviews the other address formats commonly used, and their interaction with MAC Address Randomization.

1.1. Requirements

The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL, when they appear in this document, are to be interpreted as described in [RFC2119].

2. Randomized link-layer addresses

Mobile nodes can be tracked using multiple identifiers, the most prominent being the 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 implementation of Wi-Fi use unique 48 bit MAC 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 clear text. Tracking devices can "listen to the airwaves" to find out what devices are transmitting near them.

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. MAC address randomization was successfully tried at the IETF in Honolulu in November 2014, and in several other lcation since that, as reported in [IETFMACRandom], and is also studied in [IEEE802PRSG]. MAC Address Randomization will defend against trackers that just "listen

to the airwaves," but tracking can be re-enabled if the trackers can obtain other device identifiers. We are concerned here with the use of IPv6 addresses for such tracking.

From a privacy point of view, it is clear that MAC Addresses and IPv6 addresses and DHCP identifiers shall evolve in synchrony. For example, if the MAC address changes and the IID portion of the IPv6 address stays constant, then it is really easy to correlate old and new MAC address. Conversely, if the IID changes but the MAC address remains constant, the old and new identifiers and addresses can be correlated by listening to the link's traffic.

2.1. Randomized link-layer address format

At the time of this writing, there is no standard way to construct randomized link layer addresses, but many implementations use the following algorithm for IEEE 802 48 bit MACs:

Set the the "u" (universal/local) bit to 1 (local).

Set the the "g" (individual/group) bit to 0 (individual).

Pick random values for all the other bits.

2.2. Link-layer address life time

This document makes the hypothesis that randomized link layer addresses are chosen just prior to the connection to a link. Hosts are expected to maintain the same link-layer address for the duration of the connection.

There are circumstances where a host may decide to reset its link layer address while maintaining an attachment to a link. For example, a host Ethernet interface may remain "plugged in" while the interface driver is reset to use a new MAC address. These conditions will be considered equivalent to disconnecting and then reconnecting with a new link layer address. The previously used IPv6 addresses will be discarded, and a new set of addreses will be assigned.

There are circonstances where a host may decide to reconnect to a particular link using the same link-layer address as for a previous attachment. In this case, the assignment algorithm will normally result in assigning the same IPv6 address as in the previous session, except under exceptional circumstances such as resetting the "secret key" used in [RFC7217].

3. Privacy respecting opaque identifiers

[RFC7217] specifies an algorithm that generates, for each network interface, a unique random IID per IPv6 link. The privacy properties of that algorithm depends on the specific source of the "Net_Iface" chosen by the implementer.

Most sources for the Net_IFace parameter listed in Appendix A of [RFC7217] will result in stable identifiers, independent of the link-layer address. This is useful in some deployment cases. For example, if the network interface card of a server is swapped, the specified algorithm will ensure that the server's IPv6 address do not change. However, applying the same algorithm for mobile devices enable tracking over time of a device that repeatedly visits the same location, despite attempts by the host to use different random link-layer address values.

Tracking over time is prevented if the Net_IFace parameter is set to the current link layer address. In that case, the stable addresses will have exactly the same lifetime as the link-layer identifiers. The IPv6 addresses will change whenever the link layer addresses change. Hosts that return to the same network without changing their link layer addresses will reuse the same IPv6 address. This SHOULD be the default solution for hosts implementing Link-layer Address Randomization.

Of course, this behavior could violate the statement regarding the Net_Iface parameter selection in Section 5 of [RFC7217]:

> It MUST be constant across system bootstrap sequences and other network events (e.g., bringing another interface up or down).

Although [RFC7217] isn't very specific about "other network events", it seems that it generally intends to not change for events like changing a link-layer address. For example, there is a specific statement about servers in section 5:

> ... a server-oriented operating system might prefer Net_Iface identifiers that are attached to system slots/ports, such that replacement of a NIC does not result in an IPv6 address change.

This is indeed a fine recommendation for static servers, for which [RFC7217] provides a reasonable tradeoff between stability and privacy. But for mobile hosts, the tradeoff is a bit different. We expect these mobile hosts to implement [RFC7217] as recommended by the IETF, but to also require more privacy than static servers. It turns out that a minimal update to [RFC7217] would make it suitable for these mobile hosts. The will keep the full benefits of stable

opaque identifiers when the link-layer address is stable, and the expected privacy when the link-layer address is randomized. This simple update is proposed in the next section.

3.1. Update to RFC 7217

Section 5 of [RFC7217], Net_Iface selection, is modified as follow:

> Replace "MUST" by "SHOULD" in the text:
>
> It SHOULD be constant across system bootstrap sequences and other network events (e.g., bringing another interface up or down).
>
> Immediately after that, add:
>
> It MAY change if the system administrator decides so explicitly, e.g. by implementing Link Layer Address Randomization. This can be achieved by selecting the Current Link Layer Address for Net-Iface parameter.

The following text is added to Appendix A, section A.3, Link-Layer Addresses:

> Link-Layer addresses will change dynamically in systems that implement Link Layer Address Randomization. This will cause IIDs to change whenever the Link Address changes, which is very desirable for privacy.

4. Privacy compatible Temporary Addresses

As stated in [I-D.ietf-6man-ipv6-address-generation-privacy], "a host that uses only a temporary address mitigates all four threats. Its activities may only be correlated for the lifetime a single temporary address." There is however a condition. If the lifetime of the temporary address exceeds the lifetime of the random link layer address, then correlation of successive link-layer addresses becomes possible, effectively enabling a form of tracking.

If a host uses both temporary and stable addresses, the privacy properties are those of the particular stable addresses. This is also true is a host uses temporary addresses and configure but doen't use a stable address. The address configuration will require performing duplicate address detection, generating at least a few packets on the local links. Observing this packets, an on-link attacker can correlate the link-layer address with the stable address. If the stable address includes a constant identifier, then the benefits of using random link-local addresses will be negated.

This situation is anticipated somewhat in the specification of temporary addresses. Section 3.5 of [RFC4941] specifies procedure for the regeneration of interface identifiers. The last paragraph of that section specifies:

> ... when an interface connects to a new link, a new randomized interface identifier SHOULD be generated immediately together with a new set of temporary addresses.

That condition is however not sufficient to cover the case of a device that re-connects to the same link with a new randomized link local addresses.

4.1. Update to RFC 4941

The word "Finally" should be removed from the last paragraph of section 3.5.

The following text should be added at the end of section 3.5:

> Finally, when an interface is reconfigured to use a new link-layer address, a new randomized interface identifier SHOULD be generated immediately together with a new set of temporary addresses. The previously assigned addresses SHOULD be marked as expired, not just deprecated. This reconfiguration will happen for example as a consequence of link-layer address randomization.

5. Other IPv6 Address Assigment methods

The previous sections reviewed the use of stable addresses [RFC7217] and temporary addresses [RFC4941]. Several other IPv6 address assignment methods have been defined over time. We review here these methods in light of link layer address randomization, using the same nomenclature as [I-D.ietf-6man-ipv6-address-generation-privacy]. Several IPv6 address assignment methods have been defined over time. We review here these methods in light of link layer address randomization, using the same nomenclature as [I-D.ietf-6man-ipv6-address-generation-privacy].

5.1. IEEE-identifier-based IIDs

IEEE-identifier-based IIDs could be derived from randomized link layer ID, using the algorithm specified in Appendix A of [RFC4291].

If the IIDs are constructed using the random link layer addresses, and if the random link layer addresses are constructed using the algorithm specified in Section 2.1, then the issues described in section 3 of [I-D.ietf-6man-ipv6-address-generation-privacy] are

somewhat mitigated, but many concerns remain. The correlation over time still be possible for the lifetime of the link layer address, and the location tracking will only be mitigated if link layer addresses do change with location.

In addition to the lifetime and location tracking concerns, there is also a "scope" issue with IEEE-identifier-based IIDs. The practice will export the link-layer address value to all places where the IPv6 address is used. This increase the potential "surface" for privacy attacks, and is not desirable.

There is a small probability of collision between IIDs derived from random link layer addresses and IIDs obtained through the sematically opaque, cryptographically generated, or temporary assignment methods. The "u" bit is set to global for globally assigned link layer addresses, but set to "local" for both random link layer addresses and for IIDs derived through some random process. The collision risk is however very small, and may not be a practical concern.

5.2. Static, manually configured IIDs

Because static, manually configured IIDs are stable, both correlation and location tracking are possible for the life of the address. Using randomized link-local addresses doesn't change that.

In practice, static assignment and link-layer address randomization address different scenarios. Static assignments are typically used for static hosts, while randomization is typically used for mobile hosts.

5.3. Constant, semantically opaque IIDs

This address assignment method allows correlation and location tracking because the IID is constant across IPv6 links and time. Using randomized link-local addresses doesn't change that. In fact, the constant values allow for correlation between the random link-local address and the host's identity, removing most of privacy value of random link-layer addresses.

Section 4.3 of [I-D.ietf-6man-ipv6-address-generation-privacy] addresses the general case of systems generating constant IID using the algorithms specified in [RFC4941], mentioning the implementation of this algorithm in Windows. Tests on the Windows platform show that the "constant" IIDs do in fact change if the link layer address is changed to a random value, and thus do in fact preserve the privacy value of random link-layer addresses.

5.4. DHCPv6 generation of IIDs

When using DHCPv6 in conjunction with random link layer addresses, implementers SHOULD follow the recommendations of [I-D.ietf-dhc-anonymity-profile].

5.5. Transition/co-existence technologies

Transition technologies typically embed an IPv4 address in a specifically formatted IPv6 address. Tracking over time becomes possible if the IPv4 address has a longer lifetime than the random link-layer address.

To mitigate the potential tracking issues with embedded IPv4 addresses, hosts using random link-local addresses SHOULD implement the DHCPv4 profile specified in [I-D.ietf-dhc-anonymity-profile].

6. Security Considerations

This whole document concerns the privacy and security properties of different IPv6 address generation mechanisms.

7. IANA Considerations

This draft does not require any IANA action.

8. Acknowledgments

The inspiration for this draft came from the authors of [I-D.ietf-6man-ipv6-address-generation-privacy], Alissa Cooper, Fernando Gont, and Dave Thaler. Philip Homburg and other members of the 6Man working group provided valuable comments.

9. References

9.1. 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>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <http://www.rfc-editor.org/info/rfc4291>.

[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>.

[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>.

9.2. Informative References

[I-D.ietf-6man-ipv6-address-generation-privacy] Cooper, A., Gont, F., and D. Thaler, "Privacy Considerations for IPv6 Address Generation Mechanisms", draft-ietf-6man-ipv6-address-generation-privacy-08 (work in progress), September 2015.

[I-D.ietf-dhc-anonymity-profile] Huitema, C., Mrugalski, T., and S. Krishnan, "Anonymity profile for DHCP clients", draft-ietf-dhc-anonymity-profile-08 (work in progress), February 2016.

[IEEE802PRSG] IEEE 802 EC PRSG, "IEEE 802 EC Privacy Recommendation Study Group", Dec 2015, <http://www.ieee802.org/PrivRecsg/>.

[IETFMACRandom] 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>.

Author's Address

Christian Huitema
Microsoft
Redmond, WA 98052
U.S.A.

Email: huitema@microsoft.com