

Internet Engineering Task Force (IETF)
Request for Comments: 8065
Category: Informational
ISSN: 2070-1721

D. Thaler
Microsoft
February 2017

Privacy Considerations for IPv6 Adaptation-Layer Mechanisms

Abstract

This document discusses how a number of privacy threats apply to technologies designed for IPv6 over various link-layer protocols, and it provides advice to protocol designers on how to address such threats in adaptation-layer specifications for IPv6 over such links.

Status of This Memo

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

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

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc8065>.

Copyright Notice

Copyright (c) 2017 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 (<http://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
2. Amount of Entropy Needed in Global Addresses	3
3. Potential Approaches	4
3.1. IEEE-Identifier-Based Addresses	5
3.2. Short Addresses	5
4. Recommendations	6
5. Security Considerations	7
6. Informative References	7
Author's Address	9

1. Introduction

RFC 6973 [RFC6973] discusses privacy considerations for Internet protocols, and Section 5.2 of that document covers a number of privacy-specific threats. In the context of IPv6 addresses, Section 3 of [RFC7721] provides further elaboration on the applicability of the privacy threats.

When interface identifiers (IIDs) are generated without sufficient entropy compared to the link lifetime, devices and users can become vulnerable to the various threats discussed there, including:

- o Correlation of activities over time, if the same identifier is used for traffic over period of time
- o Location tracking, if the same interface identifier is used with different prefixes as a device moves between different networks
- o Device-specific vulnerability exploitation, if the identifier helps identify a vendor or version or protocol and hence suggests what types of attacks to try
- o Address scanning, which enables all of the above attacks by off-link attackers. (On some Non-Broadcast Multi-Access (NBMA) links where all nodes aren't already privy to all on-link addresses, address scans might also be done by on-link attackers; however, in most cases, address scans are not an interesting threat from on-link attackers and thus address scans generally apply only to routable addresses.)

For example, for links that may last for years, "enough" bits of entropy means at least 46 or so bits (see Section 2 for why) in a routable address; ideally all 64 bits of the IID should be used, although historically some bits have been excluded for reasons

discussed in [RFC7421]. Link-local addresses can also be susceptible to the same privacy threats from off-link attackers, since experience shows they are often leaked by upper-layer protocols such as SMTP, SIP, or DNS.

For these reasons, [RFC8064] recommends using an address generation scheme in [RFC7217], rather than addresses generated from a fixed link-layer address.

Furthermore, to mitigate the threat of correlation of activities over time on long-lived links, [RFC4941] specifies the notion of a "temporary" address to be used for transport sessions (typically locally initiated outbound traffic to the Internet) that should not be linkable to a more permanent identifier such as a DNS name, user name, or fixed link-layer address. Indeed, the default address selection rules [RFC6724] now prefer temporary addresses by default for outgoing connections. If a device needs to simultaneously support unlinkable traffic as well as traffic that is linkable to such a stable identifier, supporting simultaneous use of multiple addresses per device is necessary.

2. Amount of Entropy Needed in Global Addresses

In terms of privacy threats discussed in [RFC7721], the one with the need for the most entropy is address scans of routable addresses. To mitigate address scans, one needs enough entropy to make the probability of a successful address probe be negligible. Typically, this is measured in the length of time it would take to have a 50% probability of getting at least one hit. Address scans often rely on sending a packet such as a TCP SYN or ICMP Echo Request, then determining whether the reply is a) an ICMP unreachable error (if no host exists with that address), b) a TCP response or ICMP Echo Reply (if a host exists), or c) none of those, in which case nothing is known for certain.

Many privacy-sensitive devices support a "stealth mode" as discussed in Section 5 of [RFC7288] or are behind a network firewall that will drop unsolicited inbound traffic (e.g., TCP SYNs, ICMP Echo Requests, etc.) and thus no TCP RST or ICMP Echo Reply will be sent. In such cases, and when the device does not listen on a well-known TCP or UDP port known to the scanner, the effectiveness of an address scan is limited by the ability to get ICMP unreachable errors, since the attacker can only infer the presence of a host based on the absence of an ICMP unreachable error.

Generation of ICMP unreachable errors is typically rate limited to 2 per second (the default in routers such as Cisco routers running IOS 12.0 or later). Such a rate results in taking about a year to completely scan 26 bits of space.

The actual math is as follows. Let 2^N be the number of devices on the subnet. Let 2^M be the size of the space to scan (i.e., M bits of entropy). Let S be the number of scan attempts. The formula for a 50% chance of getting at least one hit in S attempts is:

$$P(\text{at least one success}) = 1 - (1 - 2^N/2^M)^S = 1/2.$$

Assuming $2^M \gg S$, this simplifies to:

$$S * 2^N/2^M = 1/2, \text{ giving } S = 2^{(M-N-1)}, \text{ or } M = N + 1 + \log_2(S).$$

Using a scan rate of 2 per second, this results in the following rule of thumb:

$$\begin{aligned} \text{Bits of entropy needed} = \\ \log_2(\# \text{ devices per link}) + \log_2(\text{seconds of link lifetime}) + 2 \end{aligned}$$

For example, for a network with at most 2^{16} devices on the same long-lived link, where the average lifetime of a link is 8 years (2^{28} seconds) or less, this results in a need for at least 46 bits of entropy ($16+28+2$) so that an address scan would need to be sustained for longer than the lifetime of the link to have a 50% chance of getting a hit.

Although 46 bits of entropy may be enough to provide privacy in such cases, 59 or more bits of entropy would be needed if addresses are used to provide security against attacks such as spoofing, as CGAs [RFC3972] and HBAs [RFC5535] do, since attacks are not limited by ICMP rate limiting but by the processing power of the attacker. See those RFCs for more discussion.

If, on the other hand, the devices being scanned for respond to unsolicited inbound packets, then the address scan is not limited by the ICMP unreachable rate limit in routers, since an adversary can determine the presence of a host without them. In such cases, more bits of entropy would be needed to provide the same level of protection.

3. Potential Approaches

The table below shows the number of bits of entropy currently available in various technologies:

Technology	Reference	Bits of Entropy
802.15.4	[RFC4944]	16+ or any EUI-64
Bluetooth LE	[RFC7668]	48
DECT ULE	[DECT-ULE]	40 or any EUI-48
MS/TP	[IPv6-over-MSTP]	7 or 64
ITU-T G.9959	[RFC7428]	8
NFC	[IPv6-over-NFC]	5

Such technologies generally support either IEEE identifiers or so called "Short Addresses", or both, as link-layer addresses. We discuss each in turn.

3.1. IEEE-Identifier-Based Addresses

Some technologies allow the use of IEEE EUI-48 or EUI-64 identifiers or allow the use of an arbitrary 64-bit identifier. Using such an identifier to construct IPv6 addresses makes it easy to use the normal LOWPAN_IPHC encoding [RFC6282] with stateless compression, which allows such IPv6 addresses to be fully elided in common cases.

Global addresses with interface identifiers formed from IEEE identifiers can have insufficient entropy to mitigate address scans unless the IEEE identifier itself has sufficient entropy and enough bits of entropy are carried over into the IPv6 address to sufficiently mitigate the threats. Privacy threats other than "Correlation over time" can be mitigated using per-network randomized link-layer addresses with enough entropy compared to the link lifetime. A number of such proposals can be found at <https://mentor.ieee.org/privecsg/documents>, and Section 10.8 of [BTCorev4.1] specifies one for Bluetooth. Using routable IPv6 addresses derived from such link-layer addresses would be roughly equivalent to those specified in [RFC7217].

Correlation over time (for all addresses, not just routable addresses) can be mitigated if the link-layer address itself changes often enough, such as each time the link is established, if the link lifetime is short. For further discussion, see [RANDOM-ADDR].

Another potential concern is that of efficiency, such as avoiding Duplicate Address Detection (DAD) altogether when IPv6 addresses are based on IEEE identifiers. Appendix A of [RFC4429] provides an analysis of address-collision probability based on the number of bits of entropy. A simple web search on "duplicate MAC addresses" will show that collisions do happen with MAC addresses; thus, based on the analysis in [RFC4429], using sufficient bits of entropy in random addresses can provide greater protection against collision than using MAC addresses.

3.2. Short Addresses

A routable IPv6 address with an interface identifier formed from the combination of a "Short Address" and a set of well-known constant bits (such as padding with 0's) lacks sufficient entropy to mitigate address scanning unless the link lifetime is extremely short. Furthermore, an adversary could also use statistical methods to determine the size of the L2 address space and thereby make some inference regarding the underlying technology on a given link, and target further attacks accordingly.

When Short Addresses are desired on links that are not guaranteed to have a short enough lifetime, the mechanism for constructing an IPv6 interface identifier from a Short Address could be designed to sufficiently mitigate the problem. For example, if all nodes on a given L2 network have a shared secret (such as the key needed to get on the layer-2 network), the 64-bit IID might be generated using a one-way hash that includes (at least) the shared secret together with the Short Address. The use of such a hash would result in the IIDs being spread out among the full range of IID address space, thus mitigating address scans while still allowing full stateless compression/elision.

For long-lived links, "temporary" addresses might even be generated in the same way by (for example) also including in the hash the Version Number from the Authoritative Border Router Option (Section 4.3 of [RFC6775]), if any. This would allow changing temporary addresses whenever the Version Number is changed, even if the set of prefix or context information is unchanged.

In summary, any specification using Short Addresses should carefully construct an IID generation mechanism so as to provide sufficient entropy compared to the link lifetime.

4. Recommendations

The following are recommended for adaptation-layer specifications:

- o Security (privacy) sections should say how address scans are mitigated. An address scan might be mitigated by having a link always be short-lived, by having a large number of bits of entropy in routable addresses, or by some combination thereof. Thus, a specification should explain what the maximum lifetime of a link is in practice and show how the number of bits of entropy is sufficient given that lifetime.
- o Technologies should define a way to include sufficient bits of entropy in the IPv6 interface identifier, based on the maximum link lifetime. Specifying that randomized link-layer addresses can be used is one easy way to do so, for technologies that support such identifiers.
- o Specifications should not simply construct an IPv6 interface identifier by padding a Short Address with a set of other well-known constant bits, unless the link lifetime is guaranteed to be extremely short or the Short Address is allocated by the network (rather than being constant in the node). This also applies to link-local addresses if the same Short Address is used independent of network and is unique enough to allow location tracking.
- o Specifications should make sure that an IPv6 address can change over long periods of time. For example, the interface identifier might change each time a device connects to the network (if connections are short) or might change each day (if connections can be long). This is necessary to mitigate correlation over time.
- o If a device can roam between networks and more than a few bits of entropy exist in the IPv6 interface identifier, then make sure that the interface identifier can vary per network as the device roams. This is necessary to mitigate location tracking.

5. Security Considerations

This entire document is about security considerations and how to specify possible mitigations.

6. Informative References

- [BTCorev4.1]
Bluetooth, "Specification of the Bluetooth System",
Covered Core Package version: 4.1, December 2013,
<[https://www.bluetooth.org/DocMan/handlers/
DownloadDoc.ashx?doc_id=282159](https://www.bluetooth.org/DocMan/handlers/DownloadDoc.ashx?doc_id=282159)>.
- [DECT-ULE] Mariager, P., Petersen, J., Ed., Shelby, Z., Van de Logt,
M., and D. Barthel, "Transmission of IPv6 Packets over
DECT Ultra Low Energy", draft-ietf-6lo-dect-ule-09,
Work in Progress, December 2016.
- [IPv6-over-MSTP]
Lynn, K., Ed., Martocci, J., Neilson, C., and S.
Donaldson, "Transmission of IPv6 over MS/TP Networks",
draft-ietf-6lo-6lobac-06, Work in Progress, October 2016.
- [IPv6-over-NFC]
Choi, Y-H., Hong, Y-G., Youn, J-S., Kim, D-K., and J-H.
Choi, "Transmission of IPv6 Packets over Near Field
Communication", draft-ietf-6lo-nfc-05, Work in Progress,
October 2016.
- [RANDOM-ADDR]
Huitema, C., "Implications of Randomized Link Layers
Addresses for IPv6 Address Assignment",
draft-huitema-6man-random-addresses-03, Work in Progress,
March 2016.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)",
RFC 3972, DOI 10.17487/RFC3972, March 2005,
<<http://www.rfc-editor.org/info/rfc3972>>.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD)
for IPv6", RFC 4429, DOI 10.17487/RFC4429, April 2006,
<<http://www.rfc-editor.org/info/rfc4429>>.
- [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>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler,
"Transmission of IPv6 Packets over IEEE 802.15.4
Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007,
<<http://www.rfc-editor.org/info/rfc4944>>.

- [RFC5535] Bagnulo, M., "Hash-Based Addresses (HBA)", RFC 5535, DOI 10.17487/RFC5535, June 2009, <<http://www.rfc-editor.org/info/rfc5535>>.
- [RFC6282] Hui, J., Ed., and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <<http://www.rfc-editor.org/info/rfc6724>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<http://www.rfc-editor.org/info/rfc6775>>.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<http://www.rfc-editor.org/info/rfc6973>>.
- [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>>.
- [RFC7288] Thaler, D., "Reflections on Host Firewalls", RFC 7288, DOI 10.17487/RFC7288, June 2014, <<http://www.rfc-editor.org/info/rfc7288>>.
- [RFC7421] Carpenter, B., Ed., Chown, T., Gont, F., Jiang, S., Petrescu, A., and A. Yourtchenko, "Analysis of the 64-bit Boundary in IPv6 Addressing", RFC 7421, DOI 10.17487/RFC7421, January 2015, <<http://www.rfc-editor.org/info/rfc7421>>.
- [RFC7428] Brandt, A. and J. Buron, "Transmission of IPv6 Packets over ITU-T G.9959 Networks", RFC 7428, DOI 10.17487/RFC7428, February 2015, <<http://www.rfc-editor.org/info/rfc7428>>.

- [RFC7668] Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", RFC 7668, DOI 10.17487/RFC7668, October 2015, <<http://www.rfc-editor.org/info/rfc7668>>.
- [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>>.
- [RFC8064] Gont, F., Cooper, A., Thaler, D., and W. Liu, "Recommendation on Stable IPv6 Interface Identifiers", RFC 8064, February 2017, <<http://www.rfc-editor.org/info/rfc8064>>.

Author's Address

Dave Thaler
Microsoft
One Microsoft Way
Redmond, WA 98052
United States of America

Email: dthaler@microsoft.com

