Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 20, 2015

D. Thaler
Microsoft

February 16, 2015

Enabling Security/Privacy Addressing On 6LoWPAN Technologies
draft-thaler-6lo-privacy-addrs-00

Abstract

It is commonly assumed today that 6LowPAN header compression is incompatible (or at least inefficient) with the notion of using addresses with sufficient entropy to mitigate various security and privacy threats. This draft explores ways one might dispel that notion, and discusses how security/privacy addressing might be used on 6LoWPAN technologies without additional overhead in data packets.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

RFC 6973 [RFC6973] discusses privacy considerations for Internet protocols, and Section 5.2 in particular covers a number of privacy-specific threats. In the context of IPv6 addresses, Section 3 of [I-D.ietf-6man-ipv6-address-generation-privacy] provides further elaboration on the applicability of the privacy threats. When interface identifiers (IIDs) are generated without sufficient entropy, devices and users become vulnerable to the various threats discussed there, including correlation of activities over time, location tracking, address scanning, and device-specific vulnerability exploitation.

Interfaces identifiers formed from IEEE identifiers can have insufficient entropy 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. Typically "enough" bits of entropy means at least 46 bits (see Appendix for why); ideally all 64 bits of the IID should be used, although historically some bits have been excluded for reasons discussed in [RFC7421].

Furthermore, IEEE-identifier-based IIDs are also insufficient to prevent location tracking unless the IEEE identifier itself is different at each network location. This observation suggests that the privacy threats can be mitigated in either of two ways: either use an IPv6 address generation mechanism that is not IEEE-identifier-based, or else make sure the IEEE identifier contains at least 46 bits of entropy and is changed if a device moves to a different network. For this reason, [I-D.ietf-6man-default-iids] recommends using the address generation scheme in [RFC7217] by default, rather than IEEE-identifier-based addresses.

Furthermore, to mitigate the threat of correlation of activities over time, [RFC4941] specifies the notion of a "temporary" address to be used for sessions that should not be linkable to a more permanent identifier (such as a DNS name, user name, or stable hardware address). Such temporary addresses are appropriate for connections (typically locally-initiated outbound sessions) that an attacker cannot link to a stable identifier such as a user name or DNS name. Indeed, the default address selection rules [RFC6724] now prefer temporary addresses by default for outgoing connections. When temporary addresses are used, a new temporary address is periodically (default is 1 day in [RFC4941]) generated, which limits the threat of correlation of activies over time to that period. The address itself though may still be usable for existing long-lived connections (but not new connections) for some longer period (default is 1 week); this allows for not breaking application sessions, especially those that might be initiated shortly before a new temporary address is generated. This fact means that multiple temporary addresses can exist at the same time, one for new connections, and one or more (often up to 6, per the default periods) old ones for long-lived connections. This is in addition to any "stable" addresses that might be used for connections that are linkable to more permanent identifiers such as DNS names or user names. Whereas most threats could be mitigated if the IEEE identifier contains sufficient entropy and is different per-network, mitigating the threat of correlation of activities over time typically cannot be done using an IEEE-identifier-based-IID, since mitigating such a threat typically involves the ability to use multiple IPv6 addresses simultaneously whereas typically only one IEEE identifier can be used at a time.

Finally, allowing efficient use of addresses that are not IEEE-identifier-based also has additional security benefits not specific to privacy. For example, addresses such as Cryptographically Generated Addresses (CGAs) [RFC3972] and Hash-Based Addresses (HBAs) [RFC5535] can be used in security protocols such as Secure Neighbor Discovery (SeND) [RFC6496], IPsec, etc. Such techniques rely on having around 59 or more bits of entropy in the address to provide sufficient cryptographic protection.

RFC 6775 [RFC6775] already allows for the use of non-IEEE-identifier-based addresses, such as those provided by DHCPv6 [RFC3315]. There has been some concern, however, that such approaches necessarily interfere with efficient header compression for IPv6 (e.g., over IEEE 802.15.4-based networks [RFC6282]), as it is important to keep data packets small on 6LoWPAN networks.

Another potential concern is that of efficiency, such as avoiding DAD all together when IPv6 addresses are IEEE-identifier-based. 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, and thus based on the analysis in [RFC4429], using sufficient bits of entropy in non-IEEE-identifier-based addresses can provide greater protection against collision than using MAC addresses.

The remainder of this document explores how one might use addresses with sufficient entropy on 6LoWPAN networks while avoiding extra overhead.

2. Terminology

This document uses the terminology defined in Section 3 of [RFC6973], including terms such as "(un)linkability" and "anonymity set".

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].

3. Compression Details

The LOWPAN_IPHC encoding format specified in Section 3.1 of RFC 6282 [RFC6282] defines a method for deriving IIDs from the link-layer source and/or destination addresses in the encapsulation header. Unicast IPv6 addresses may be compressed to 64, 16, or 0 bits in the encoded IPv6 header.

3.1. Use of IEEE-Identifier-Based Addresses

As noted earlier, some threats could be mitigated using per-network "randomized" IEEE identifiers with 46 or more bits of entropy. 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 IPv6 addresses derived from such IEEE identifiers would be roughly equivalent to those specified in [RFC7217].

Such addresses would be encoded as usual using the LOWPAN_IPHC encoding format. For example, if the source and destination addresses are both on-link and derived from the IEEE identifier in the encapsulating header:

- SAC (Source Address Compression) is set to 0 to indicate stateless compression.

- SAM (Source Address Mode) is set to 11 to indicate the address is fully elided and can be computed from the encapsulating header.

- DAC (Destination Address Compression) is set to 0 to indicate stateless compression.

- DAM (Destination Address Mode) is set to 11 to indicate the address is fully elided and can be computed from the encapsulating header.

3.2. Use of 16-Bit Short Addresses

An IPv6 address formed (per Section 6 of [RFC4944]) from an 16-bit identifier such as an IEEE 802.15.4 16-bit short address does not provide sufficient entropy to fully mitigate address scanning, as the size of the address scan search space depends on the entropy in the IID, and only 15 bits are available for unicast addresses. An adversary could also use statisical methods to determine the size of the L2 address space and thereby make some inference regarding the underlying technology being IEEE 802.15.4 on a given link. As such, this address generation mechanism SHOULD NOT be used on networks where privacy threats may be an issue, such as any networks that have Internet connectivity.

It might be possible to construct IPv6 addresses from 16-bit short addresses using an alternate mechanism that mitigates address scans, if all nodes on a given L2 network have a shared secret (such as the key needed to get on the layer-2 network) and generate the IID by using a one-way 64-bit hash of 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.

"Temporary" addresses could possibly be generated in the same way by also including in the hash the Version Number from the Authoritative Border Router Option (ABDO) 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). Such a scheme would likely require using the Context Identifier (CID) to distinguish between non-temporary addresses, "current" temporary addresses, and "past" temporary addresses based on a previous Version Number.

Specifying further details of such a scheme is left for future versions of this draft, if there is interest.

3.3. Use of Non-IEEE-Identifier-Based Addresses

Unicast addresses that are not IEEE-identifier based could be compressed to 0 bits as follows, using stateful context-based compression where the entire IPv6 address including the IID (as opposed to only the IPv6 prefix) are covered by context information. It is also worth pointing out that this same scheme would also allow compressing DHCPv6-assigned addresses even in networks where privacy is not a primary concern, thus potentially providing efficiency benefits in addition to privacy and security ones. Furthermore, unlike stateless compression, stateful context-based compression could also allow compressing addresses of nodes outside the local network (i.e., where the IEEE identifier in the encapsulating header is that of a router rather than the peer, and the peer’s address does not have a prefix in the local network) and hence can provide greater savings in such cases.

3.3.1. Source Address Compression

SAC (Source Address Compression) MUST be set to 1 to indicate stateful context-based compression.

SAM (Source Address Mode) MUST be set to 11 to indicate that the address is fully elided.

3.3.2. Destination Address Compression

DAC (Destination Address Compression) MUST be set to 1 to indicate stateful context-based compression.

DAM (Destination Address Mode) MUST be set to 11 to indicate that the address is fully elided.

3.3.3. Context Identifier

When non-IEEE-identifier-based addresses are used as described in this document, each address MUST be associated with a separate context. That is, the "prefix" associated with a context MUST be the full 128 bits of the IPv6 address.

LOWPAN_IPHC supports up to 16 source address contexts and 16 destination address contexts, allowing for simultaneous use of up to

16 source addresses and 16 destination addresses that are not IEEE-identifier-based. Context 0 is the default context if the CID (Context Identifier Extension) octet is absent, and other values require the CID to be present. As such, the address most commonly used (typically either the stable non-temporary address, or the currently preferred temporary address) could be assigned to context 0 so that the presence of the CID octet is minimized.

3.3.4. Context State

As specified in [RFC6775], context state is distributed by routers and is shared across a LoWPAN. This means that the use of CIDs described above would only support compression of 16 source and destination addresses across the entire LoWPAN. However, Section 8 of [RFC6775] explicitly allows for such context dissemination to be substituted by alternatives defined in other specifications. We now describe such a substitute that would allow header compression with up to 16 source addresses and 16 destination addresses *per node*.

First, a context entry is defined to be indexed by a { link-layer address, CID } tuple, rather than just a CID. Second, each node is responsible for generating and disseminating the CIDs for its own IPv6 addresses.

Thus, each Neighbor Cache Entry (NCE) in a router conceptually contains the CID of the neighbor’s address, used when compressing packets sent to it.

3.3.5. Context Distribution

To disseminate CID information from a host to a router, the Address Registration Option (ARO) defined in Section 4.1 of [RFC6775] can be extended to include the CID by using 5 of the 24 Reserved bits (one for a flag to denote a CID is present, and 4 for the CID). For distribution in a multihop network, the Duplicate Address Request (DAR) and Duplicate Address Confirmation (DAC) messages can be similarly extended to include the CID in currently Reserved bits.

To disseminate CID information from a router to a host, Section 4.2 of [RFC6775] defines the 6LoWPAN Context Option (6CO) for use in Router Discovery. If a router sees that a host is sending packets without compressing a source or destination address, the router could send it an updated 6CO with a CID for that address as the context prefix, to allow compression of subsequent packets. Since each non-IEEE-identifier-based address requires its own context, the Context Length field MUST be set to 128 in the 6CO containing such context information. Note that the CID in a 6CO for another address within the 6LoWPAN is still generated by the router (since it is specific to the router’s link-layer address as used by the host to which the 6CO is sent); it is not the same value as the CID generated by the destination node itself, which CID is used by its router when forwarding a packet to it. Thus a router is responsible for updating CIDs in packets it forwards, just as it updates the link-layer source and destination addresses in the encapsulating header.

Specifying further details of such a scheme is left for future versions of this draft, if there is interest.

3.3.6. Negotiation

To negotiate using the substitute mechanisms above, rather than the default mechanisms specified in [RFC6775], the 6LoWPAN Capability Indication Option (6CIO) could be used as allowed for in Section 3.4 of [RFC7400] by assigning one of the "6LoWPAN capability Bits" for this purpose.

3.3.7. Discussion of Tradeoffs

This proposal decentralizes a portion of context generation and distribution to include simple nodes. In many 6LoWPAN scenarios, as much as possible is offloaded to router nodes precisely because end nodes are so limited. Until context info is learned for a given destination address, a node is not able to compress it. Compression would kick in after the context info is known. After context info is learned, the 4-bit CID must be stored for the destination address. As such, using this scheme requires a slight amount of overhead in the initial packet(s) but no additional overhead afterwards, and it requires no additional memory overhead initially, but a slight amount of additional memory overhead after context is learned.

In the rare case that a simple node needs to simultaneously communicate with more than 16 other non-IEEE-identifier-based destination addresses, at most 16 of them will be able to be compressed, and the others will have additional packet overhead.

4. IANA Considerations

The approach described in Section 3.3 would require IANA to allocate a bit in the "6LoWPAN capability Bits" subregistry for this purpose.

5. Security Considerations

This entire document is about security considerations and possible mitigations.

6. Acknowledgements

Thanks to Fernando Gont, Christian Huitema, and Gabriel Montenegro for discussion on the ideas described in this draft.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, September 2007.

[RFC6282] Hui, J. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, September 2011.

[RFC6775] Shelby, Z., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, November 2012.

[RFC7400] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 7400, November 2014.

7.2. Informative References

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.

[RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, April 2006.

[RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, September 2007.

[RFC5535] Bagnulo, M., "Hash-Based Addresses (HBA)", RFC 5535, June 2009.

[RFC6496] Krishnan, S., Laganier, J., Bonola, M., and A. Garcia-Martinez, "Secure Proxy ND Support for SEcure Neighbor Discovery (SEND)", RFC 6496, February 2012.

[RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, July 2013.

[RFC7136] Carpenter, B. and S. Jiang, "Significance of IPv6 Interface Identifiers", RFC 7136, February 2014.

[RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, April 2014.

[RFC7288] Thaler, D., "Reflections on Host Firewalls", RFC 7288, June 2014.

[RFC7421] Carpenter, B., Chown, T., Gont, F., Jiang, S., Petrescu, A., and A. Yourtchenko, "Analysis of the 64-bit Boundary in IPv6 Addressing", RFC 7421, January 2015.

[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-03 (work in progress), January 2015.

[I-D.ietf-6man-default-iids]
Gont, F., Cooper, A., Thaler, D., and W. Will, "Recommendation on Stable IPv6 Interface Identifiers", draft-ietf-6man-default-iids-02 (work in progress), January 2015.

[BTCorev4.1]
Bluetooth Special Interest Group, "Bluetooth Core Specification Version 4.1", December 2013, <https://www.bluetooth.org/DocMan/handlers/DownloadDoc.ashx?doc_id=282159>.

Appendix A. Amount of Entropy Needed

In terms of privacy threats discussed in [I-D.ietf-6man-ipv6-address-generation-privacy], the one with the need for the most entropy is address scans. 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, and determining whether the reply is an ICMP unreachable errors (if no host exists) or TCP response or ICMP Echo Reply (if a host exists), or neither in which case nothing is known for certain.

Many privacy-sensitive devices support a "stealth mode" as discussed in Section 5 of [RFC7288] whereby they will not send a TCP RST or ICMP Echo Reply. In such cases, and when the device does not listen on a well-known TCP 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 absense 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. For a network with at most 2^16 devices on the same subnet, and the average lifetime of a device being 16 (2^4) years or less, this results in a need for at least 46 bits of entropy (16+26+4) so that a address scan would need to be sustained for longer than the lifetime of devices to have a 50% chance of getting a hit.

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 is: P(at least one success) = 1 - (1 - 2^N/2^M)^S = 1/2. Assuming 2^M >> S, this simplifies to: S * 2^N/2^M = 1/2, giving S = 2^(M-N) / 2, or M = N + log_2(2S).

Although 46 bits of entropy may be enough to provide privacy in such cases, 59 or more bits of entropy are 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 do not implement a "stealth mode", but respond with TCP RST or ICMP Echo Reply packets, then the address scan is not limited by the ICMP unreachable rate limit in routers, since the attacker 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.

Author’s Address

Dave Thaler
Microsoft
One Microsoft Way
Redmond, WA 98052
USA

Email: dthaler@microsoft.com