Significance of IPv6 Interface Identifiers

Abstract

   The IPv6 addressing architecture includes a unicast interface
   identifier that is used in the creation of many IPv6 addresses.
   Interface identifiers are formed by a variety of methods. This
   document clarifies that the bits in an interface identifier have no
   meaning and that the entire identifier should be treated as an opaque
   value. In particular, RFC 4291 defines a method by which the
   Universal and Group bits of an IEEE link-layer address are mapped
   into an IPv6 unicast interface identifier. This document clarifies
   that those two bits are significant only in the process of deriving
   interface identifiers from an IEEE link-layer address, and it updates
   RFC 4291 accordingly.

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/rfc7136.

Copyright Notice

Table of Contents

1. Introduction

   IPv6 unicast addresses consist of a prefix followed by an Interface
   Identifier (IID). The IID is supposed to be unique on the links
   reached by routing to that prefix, giving an IPv6 address that is
   unique within the applicable scope (link local or global). According
   to the IPv6 addressing architecture [RFC4291], when a 64-bit IPv6
   unicast IID is formed on the basis of an IEEE EUI-64 address, usually
   itself expanded from a 48-bit MAC address, a particular format must
   be used:

      For all unicast addresses, except those that start with the binary
      value 000, Interface IDs are required to be 64 bits long and to be
      constructed in Modified EUI-64 format.

   Thus, the specification assumes that the normal case is to transform
   an Ethernet-style address into an IID, but, in practice, there are
   various methods of forming such an IID.

   The Modified EUI-64 format preserves the information provided by two
   particular bits in the MAC address:

   o The "u/l" bit in a MAC address [IEEE802] is set to 0 to indicate
      universal scope (implying uniqueness) or to 1 to indicate local
      scope (without implying uniqueness). In an IID formed from a MAC
      address, this bit is simply known as the "u" bit and its value is
      inverted, i.e., 1 for universal scope and 0 for local scope.
      According to [RFC4291] and [RFC7042], the reason for this was to
      make it easier for network operators to manually configure
      local-scope IIDs.

      In an IID, this bit is in position 6, i.e., position 70 in the
      complete IPv6 address (when counting from 0).

   o The "i/g" bit in a MAC address is set to 1 to indicate group
      addressing (link-layer multicast). The value of this bit is
      preserved in an IID, where it is known as the "g" bit.

      In an IID, this bit is in position 7, i.e., position 71 in the
      complete IPv6 address (when counting from 0).

   This document discusses problems observed with the "u" and "g" bits
   as a result of the above requirements and the fact that various other
   methods of forming an IID have been defined independently of the
   method described in Appendix A of RFC 4291. It then discusses the
   usefulness of these two bits and the significance of the bits in an
   IID in general. Finally, it updates RFC 4291 accordingly.

1.1. Terminology

   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. Problem Statement

   In addition to IIDs formed from IEEE EUI-64 addresses, various new
   forms of IIDs have been defined, including temporary addresses
   [RFC4941], Cryptographically Generated Addresses (CGAs) [RFC3972]
   [RFC4982], Hash-Based Addresses (HBAs) [RFC5535], and ISATAP
   addresses [RFC5214]. Other methods have been proposed, such as
   stable privacy addresses [IID-SLAAC] and mapped addresses for 4rd
   [SOFTWR-4RD]. In each case, the question of how to set the "u" and
   "g" bits has to be decided. For example, RFC 3972 specifies that
   they are both zero in CGAs, and RFC 4982 describes them as if they
   were reserved bits. The same applies to HBAs. On the other hand,
   RFC 4941 specifies that "u" must be zero but leaves "g" variable.

The NAT64 addressing format [RFC6052] sets the whole byte containing "u" and "g" to zero.

Another case where the "u" and "g" bits are specified is in the Reserved IPv6 Subnet Anycast Address format [RFC2526], which states that "for interface identifiers in EUI-64 format, the universal/local bit in the interface identifier MUST be set to 0" (i.e., local) and the "g" bit is required to be set to 1. However, the text neither states nor implies any semantics for these bits in anycast addresses.

A common operational practice for well-known servers is to manually assign a small number as the IID, in which case "u" and "g" are both zero.

These cases illustrate that the statement quoted above from RFC 4291 requiring "Modified EUI-64 format" is inapplicable when applied to forms of IID that are not in fact based on an underlying EUI-64 address. In practice, the IETF has chosen to assign some 64-bit IIDs that have nothing to do with EUI-64.

A particular case is that of /127 prefixes for point-to-point links between routers, as standardised by [RFC6164]. The addresses on these links are undoubtedly global unicast addresses, but they do not have a 64-bit IID. The bits in the positions named "u" and "g" in such an IID have no special significance and their values are not specified.

Each time a new IID format is proposed, the question arises whether these bits have any meaning. Section 2.2.1 of [RFC7042] discusses the mechanics of the bit allocations but does not explain the purpose or usefulness of these bits in an IID. There is an IANA registry for reserved IID values [RFC5453], but again there is no explanation of the purpose of the "u" and "g" bits.

There was a presumption when IPv6 was designed and the IID format was first specified that a universally unique IID might prove to be very useful, for example to contribute to solving the multihoming problem. Indeed, the addressing architecture [RFC4291] states this explicitly:

   The use of the universal/local bit in the Modified EUI-64 format
   identifier is to allow development of future technology that can
   take advantage of interface identifiers with universal scope.

However, so far, this has not proved to be the case. Also, there is evidence from the field that MAC addresses with universal scope are sometimes assigned to multiple MAC interfaces. There are recurrent reports of manufacturers assigning the same MAC address to multiple devices, and significant reuse of the same virtual MAC address is

reported in virtual machine environments. Once transformed into IID
format (with "u" = 1), these identifiers would purport to be
universally unique but would in fact be ambiguous. This has no known
harmful effect as long as the replicated MAC addresses and IIDs are
used on different layer 2 links. If they are used on the same link,
of course there will be a problem, very likely interfering with
link-layer transmission. If not, the problem will be detected by
duplicate address detection [RFC4862] [RFC6775], but such an error
can usually only be resolved by human intervention.

The conclusion from this is that the "u" bit is not a reliable
indicator of universal uniqueness.

We note that Identifier-Locator Network Protocol (ILNP), a
multihoming solution that might be expected to benefit from
universally unique IIDs in modified EUI-64 format, does not in fact
rely on them. ILNP uses its own format defined as a Node Identifier
[RFC6741]. ILNP has the constraint that a given Node Identifier must
be unique within the context of a given Locator (i.e., within a
single given IPv6 subnetwork). As we have just shown, the state of
the "u" bit does not in any way guarantee such uniqueness, but
duplicate address detection is available.

Thus, we can conclude that the value of the "u" bit in IIDs has no
particular meaning. In the case of an IID created from a MAC address
according to RFC 4291, its value is determined by the MAC address,
but that is all.

An IPv6 IID should not be created from a MAC group address, so the
"g" bit will normally be zero. But, this value also has no
particular meaning. Additionally, the "u" and the "g" bits are both
meaningless in the format of an IPv6 multicast group ID [RFC3306]
[RFC3307].

None of the above implies that there is a problem with using the "u"
and "g" bits in MAC addresses as part of the process of generating
IIDs from MAC addresses, or with specifying their values in other
methods of generating IIDs. What it does imply is that after an IID
is generated by any method, no reliable deductions can be made from
the state of the "u" and "g" bits; in other words, these bits have no
useful semantics in an IID.

Once this is recognised, we can avoid the problematic confusion
caused by these bits each time that a new form of IID is proposed.

3. Usefulness of the U and G Bits

   Given that the "u" and "g" bits do not have a reliable meaning in an
   IID, it is relevant to consider what usefulness they do have.

   If an IID is known or guessed to have been created according to
   [RFC4291], it could be transformed back into a MAC address. This can
   be very helpful during operational fault diagnosis. For that reason,
   mapping the IEEE "u" and "g" bits into the IID has operational
   usefulness. However, it should be stressed that an IID with "u" = 1
   and "g" = 0 might not be formed from a MAC address; on the contrary,
   it might equally result from another method. With other methods,
   there is no reverse transformation available.

   Given that the values of the "u" and "g" bits in an IID have no
   particular meaning, new methods of IID formation are at liberty to
   use them as they wish, for example, as additional pseudo-random bits
   to reduce the chances of duplicate IIDs.

4. The Role of Duplicate Address Detection

   As mentioned above, Duplicate Address Detection (DAD) [RFC4862] is
   able to detect any case where a collision of two IIDs on the same
   link leads to a duplicated IPv6 address. The scope of DAD may be
   extended to a set of links by a DAD proxy [RFC6957] or by Neighbor
   Discovery Optimization [RFC6775]. Since DAD is mandatory for all
   nodes, there will be almost no case in which an IID collision,
   however unlikely it may be, is not detected. It is out of scope of
   most existing specifications to define the recovery action after a
   DAD failure, which is an implementation issue. If a manually created
   IID, or an IID derived from a MAC address according to RFC 4291,
   leads to a DAD failure, human intervention will most likely be
   required. However, as mentioned above, some methods of IID formation
   might produce IID values with "u" = 1 and "g" = 0 that are not based
   on a MAC address. With very low probability, such a value might
   collide with an IID based on a MAC address.

   As stated in RFC 4862:

      On the other hand, if the duplicate link-local address is not
      formed from an interface identifier based on the hardware address,
      which is supposed to be uniquely assigned, IP operation on the
      interface MAY be continued.

   Continued operation is only possible if a new IID is created. The
   best procedure to follow for this will depend on the IID formation
   method in use. For example, if an IID is formed by a pseudo-random
   process, that process could simply be repeated.

5. Clarification of Specifications

   This section describes clarifications to the IPv6 specifications that
   result from the above discussion.

   The EUI-64 to IID transformation defined in the IPv6 addressing
   architecture [RFC4291] MUST be used for all cases where an IPv6 IID
   is derived from an IEEE MAC or EUI-64 address. With any other form
   of link-layer address, an equivalent transformation SHOULD be used.

   Specifications of other forms of 64-bit IIDs MUST specify how all 64
   bits are set, but a generic semantic meaning for the "u" and "g" bits
   MUST NOT be defined. However, the method of generating IIDs for
   specific link types MAY define some local significance for certain
   bits.

   In all cases, the bits in an IID have no generic semantics; in other
   words, they have opaque values. In fact, the whole IID value MUST be
   viewed as an opaque bit string by third parties, except possibly in
   the local context.

   The following statement in Section 2.5.1 of the IPv6 addressing
   architecture [RFC4291]:

      For all unicast addresses, except those that start with the binary
      value 000, Interface IDs are required to be 64 bits long and to be
      constructed in Modified EUI-64 format.

   is replaced by:

      For all unicast addresses, except those that start with the binary
      value 000, Interface IDs are required to be 64 bits long. If
      derived from an IEEE MAC-layer address, they must be constructed
      in Modified EUI-64 format.

   The following statement in Section 2.5.1 of the IPv6 addressing
   architecture [RFC4291] is obsoleted:

      The use of the universal/local bit in the Modified EUI-64 format
      identifier is to allow development of future technology that can
      take advantage of interface identifiers with universal scope.

   As far as is known, no existing implementation will be affected by
   these changes. The benefit is that future design discussions are
   simplified.

6. Security Considerations

   No new security exposures or issues are raised by this document.

   In some contexts, unpredictable IID values are considered beneficial
   to enhance privacy and defeat scanning attacks. The recognition that
   the IID value should be regarded as an opaque bit string is
   consistent with methods of IID formation that result in
   unpredictable, pseudo-random values.

7. IANA Considerations

   This document requests no immediate action by IANA. However, the
   following should be noted when considering any future proposed
   addition to the registry of reserved IID values, which requires
   Standards Action [RFC5226] according to [RFC5453].

   Full deployment of a new reserved IID value would require updates to
   IID generation code in every deployed IPv6 stack, so the technical
   justification for such a Standards Action would need to be extremely
   strong.

   The preceding sentence and a reference to this document have been
   added to the "Reserved IPv6 Interface Identifiers" registry.

8. Acknowledgements

   Valuable comments were received from Ran Atkinson, Remi Despres,
   Ralph Droms, Fernando Gont, Eric Gray, Brian Haberman, Joel Halpern,
   Bob Hinden, Christian Huitema, Ray Hunter, Tatuya Jinmei, Roger
   Jorgensen, Mark Smith, Bernie Volz, and other participants in the
   6MAN working group.

   Brian Carpenter was a visitor at the Computer Laboratory, Cambridge
   University during part of this work.

9. References

9.1. Normative References

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

   [RFC4291]   Hinden, R. and S. Deering, "IP Version 6 Addressing
               Architecture", RFC 4291, February 2006.

   [RFC4862]   Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless
               Address Autoconfiguration", RFC 4862, September 2007.

   [RFC5453] Krishnan, S., "Reserved IPv6 Interface Identifiers", RFC
              5453, February 2009.

   [RFC7042] Eastlake, D. and J. Abley, "IANA Considerations and IETF
              Protocol and Documentation Usage for IEEE 802 Parameters",
              BCP 141, RFC 7042, October 2013.

9.2. Informative References

   [IEEE802] "IEEE Standard for Local and Metropolitan Area Networks:
              Overview and Architecture", IEEE Std 802-2001 (R2007),
              2007.

   [IID-SLAAC]
              Gont, F., "A method for Generating Stable Privacy-Enhanced
              Addresses with IPv6 Stateless Address Autoconfiguration
              (SLAAC)", Work in Progress, March 2012.

   [RFC2526] Johnson, D. and S. Deering, "Reserved IPv6 Subnet Anycast
              Addresses", RFC 2526, March 1999.

   [RFC3306] Haberman, B. and D. Thaler, "Unicast-Prefix-based IPv6
              Multicast Addresses", RFC 3306, August 2002.

   [RFC3307] Haberman, B., "Allocation Guidelines for IPv6 Multicast
              Addresses", RFC 3307, August 2002.

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

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

   [RFC4982] Bagnulo, M. and J. Arkko, "Support for Multiple Hash
              Algorithms in Cryptographically Generated Addresses
              (CGAs)", RFC 4982, July 2007.

   [RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site
              Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214,
              March 2008.

   [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an
              IANA Considerations Section in RFCs", BCP 26, RFC 5226,
              May 2008.

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

   [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X.
              Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052,
              October 2010.

   [RFC6164] Kohno, M., Nitzan, B., Bush, R., Matsuzaki, Y., Colitti,
              L., and T. Narten, "Using 127-Bit IPv6 Prefixes on Inter-
              Router Links", RFC 6164, April 2011.

   [RFC6741] Atkinson,, RJ., "Identifier-Locator Network Protocol
              (ILNP) Engineering Considerations", RFC 6741, November
              2012.

   [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.

   [RFC6957] Costa, F., Combes, J-M., Pougnard, X., and H. Li,
              "Duplicate Address Detection Proxy", RFC 6957, June 2013.

   [SOFTWR-4RD]
              Despres, R., Jiang, S., Penno, R., Lee, Y., Chen, G., and
              M. Chen, "IPv4 Residual Deployment via IPv6 - a Stateless
              Solution (4rd)", Work in Progress, October 2013.

Authors' Addresses

   Brian Carpenter
   Department of Computer Science
   University of Auckland
   PB 92019
   Auckland 1142
   New Zealand

   EMail: brian.e.carpenter@gmail.com

   Sheng Jiang
   Huawei Technologies Co., Ltd
   Q14, Huawei Campus
   No.156 Beiqing Road
   Hai-Dian District, Beijing 100095
   P.R. China

   EMail: jiangsheng@huawei.com