

6MAN
Internet-Draft
Updates: [4291](#) (if approved)
Intended status: Standards Track
Expires: November 26, 2013

B. E. Carpenter
Univ. of Auckland
S. Jiang
Huawei Technologies Co., Ltd
May 25, 2013

Significance of IPv6 Interface Identifiers
draft-ietf-6man-ug-01

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 generic meaning and that the 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 bits apply only to interface identifiers that are derived from an IEEE link-layer address. It updates [RFC 4291](#) accordingly.

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 November 26, 2013.

Copyright Notice

Copyright (c) 2013 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
1.1.	Terminology	3
2.	Problem statement	3
3.	Usefulness of the U and G Bits	5
4.	The Role of Duplicate Address Detection	6
5.	Clarification of Specifications	6
6.	Security Considerations	7
7.	IANA Considerations	8
8.	Acknowledgements	8
9.	Change log [RFC Editor: Please remove]	8
10.	References	9
10.1.	Normative References	9
10.2.	Informative References	9
	Authors' Addresses	10

[1.](#) Introduction

IPv6 unicast addresses consist of a subnet prefix followed by an Interface Identifier (IID), the latter supposedly unique on the links reached by routing to that prefix. 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 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 interface identifier.

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

- o The "u" bit in an IEEE address is set to 0 to indicate universal scope (implying uniqueness) or to 1 to indicate local scope (without implying uniqueness). In an IID this bit is inverted,

i.e., 1 for universal scope and 0 for local scope. According to [RFC 4291](#) and [[RFC5342](#)], 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.

- o The "g" bit in an IEEE address is set to 1 to indicate group addressing (link-layer multicast). The value of this bit is preserved in an IID.

In an IID, this bit is in position 7, i.e., position 71 in the complete IPv6 address.

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, quite 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 IID have been defined or proposed, such as temporary addresses [[RFC4941](#)], Cryptographically Generated Addresses (CGAs) [[RFC3972](#)], Hash-Based Addresses (HBAs) [[RFC5535](#)], stable privacy addresses [[I-D.ietf-6man-stable-privacy-addresses](#)], or mapped addresses for 4rd [[I-D.ietf-softwire-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 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 requires that "g" bit 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 rather meaningless 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 RFC 5342](#) 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, this has not so far proved to be the case. Also, there is evidence from the field that IEEE MAC addresses with "u" = 0 are sometime incorrectly assigned to multiple MAC interfaces. Firstly, there are recurrent reports of manufacturers assigning the same MAC address to multiple devices. Secondly, significant re-use 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, to be detected by duplicate address detection [\[RFC4862\]](#), but such a problem 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 [RFC 4291](#), 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 "u" = "g" = 0 does not prove that an IID was formed from a MAC address; on the contrary, it might equally result from another method. With other methods, there is no reverse transformation available.

To the extent that each method of IID creation specifies the values of the "u" and "g" bits, and that each new method is carefully designed in the light of its predecessors, these bits tend 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 [[I-D.ietf-6man-dad-proxy](#)]. Since DAD is mandatory for all nodes, there will be 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. The best procedure to follow will depend on the IID formation method in use. For example, if an IID is formed by some pseudo-random process, that process could simply be repeated. 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.

There is one case in [RFC 4862](#) that requires additional consideration:

"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."

However, as mentioned above, some methods of IID formation might produce IID values with "u" = "g" = 0 that are not based on a MAC (hardware) address. With very low probability, such a value might collide with an IID based on a MAC address. There is no algorithm for determining whether this case has arisen, rather than a genuine MAC address collision. Implementers should carefully consider the consequences of continuing IPv6 operation on the interface in this unlikely situation.

5. Clarification of Specifications

This section describes clarifications to the IPv6 specifications that result from the above discussion. Their aim is to reduce confusion while retaining the useful aspects of the "u" and "g" bits in IIDs.

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 IID MUST specify how all 64 bits are set, but need not treat the "u" and "g" bits in any special way. A general semantic meaning for these 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 general 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.

7. IANA Considerations

This document requests no immediate action by IANA. However, the following should be noted when considering future proposed additions to the registry of reserved IID values, which requires Standards Action 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.

OPEN ISSUE: Alternatively, we could decide to close the reserved IID registry completely (which would also mean formally updating [RFC 5453](#)). If we choose this approach, the following point can be deleted. Comments welcome.

A reserved IID, or a range of reserved IIDs, will most likely specify values for both "u" and "g", since they are among the high-order bits. At the present time, none of the known methods of generating IIDs will generate "u" = "g" = 1. Reserved IIDs with "u" = "g" = 1 are therefore unlikely to collide with automatically generated IIDs.

8. Acknowledgements

Valuable comments were received from Ran Atkinson, Remi Despres, Fernando Gont, Brian Haberman, Joel Halpern, Bob Hinden, Christian Huitema, Ray Hunter, Mark Smith, and other participants in the 6MAN working group.

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

This document was produced using the xml2rfc tool [[RFC2629](#)].

9. Change log [RFC Editor: Please remove]

[draft-ietf-6man-ug-01](#): emphasised "opaque" nature of IID, added text about DAD failures, expanded IANA considerations, 2013-05-25.

[draft-ietf-6man-ug-00](#): first WG version, text clarified, added possibility of link-local significance, 2013-03-28.

[draft-carpenter-6man-ug-01](#): numerous clarifications following WG comments, discussed DAD, added new section on the usefulness of the u/g bits, expanded IANA considerations, set intended status, 2013-02-21.

[draft-carpenter-6man-ug-00](#): original version, 2013-01-31.

10. References

10.1. 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.
- [RFC5342] Eastlake, D., "IANA Considerations and IETF Protocol Usage for IEEE 802 Parameters", [BCP 141](#), [RFC 5342](#), September 2008.
- [RFC5453] Krishnan, S., "Reserved IPv6 Interface Identifiers", [RFC 5453](#), February 2009.

10.2. Informative References

- [I-D.ietf-6man-dad-proxy] Costa, F., Combes, J., Pournard, X., and L. Hongyu, "Duplicate Address Detection Proxy", [draft-ietf-6man-dad-proxy-07](#) (work in progress), April 2013.
- [I-D.ietf-6man-stable-privacy-addresses] Gont, F., "A method for Generating Stable Privacy-Enhanced Addresses with IPv6 Stateless Address Autoconfiguration (SLAAC)", [draft-ietf-6man-stable-privacy-addresses-07](#) (work in progress), May 2013.
- [I-D.ietf-softwire-4rd] Despres, R., Jiang, S., Penno, R., Lee, Y., Chen, G., and M. Chen, "IPv4 Residual Deployment via IPv6 - a Stateless Solution (4rd)", [draft-ietf-softwire-4rd-05](#) (work in progress), April 2013.
- [RFC2526] Johnson, D. and S. Deering, "Reserved IPv6 Subnet Anycast Addresses", [RFC 2526](#), March 1999.
- [RFC2629] Rose, M.T., "Writing I-Ds and RFCs using XML", [RFC 2629](#), June 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.
- [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, R.J., "Identifier-Locator Network Protocol (ILNP) Engineering Considerations", [RFC 6741](#), November 2012.

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

