

6MAN  
Internet-Draft  
Updates: 4291 (if approved)  
Intended status: Standards Track  
Expires: August 25, 2013

B. Carpenter  
Univ. of Auckland  
S. Jiang  
Huawei Technologies Co., Ltd  
February 21, 2013

The U and G bits in IPv6 Interface Identifiers  
draft-carpenter-6man-ug-01

Abstract

The IPv6 addressing architecture 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 the status of those bits for interface identifiers that are not derived from an IEEE link-layer address, and 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 August 25, 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 . . . . . 3
  - 1.1. Terminology . . . . . 3
- 2. Problem statement . . . . . 3
- 3. Usefulness of the U and G Bits . . . . . 5
- 4. Clarification of Specifications . . . . . 6
- 5. Security Considerations . . . . . 6
- 6. IANA Considerations . . . . . 7
- 7. Acknowledgements . . . . . 7
- 8. Change log [RFC Editor: Please remove] . . . . . 7
- 9. References . . . . . 7
  - 9.1. Normative References . . . . . 7
  - 9.2. Informative References . . . . . 8
- Authors' Addresses . . . . . 9

## 1. Introduction

According to the IPv6 addressing architecture [RFC4291], when a 64-bit IPv6 unicast Interface Identifier (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."

The specification assumes that that the normal case is to transform an Ethernet-style address into an IID, preserving the information provided by two bits in particular:

- 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 [RFC5342], the reason for this was "to make it easier for network operators to type in local-scope identifiers".
- 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.

This document discusses problems observed with the "u" and "g" bits as a result of the above requirements. It then discusses the usefulness of the two bits, and 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.

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.

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 does have the constraint that Node Identifiers must be unique within a given site, but as we have just shown, the state of the "u" bit does not in any way guarantee this.

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. 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. However, the resulting "u" and "g" bits in an IID have no semantics; in other words, they have opaque values. In fact, the whole IID should be viewed as opaque by third parties.

Specifications of other forms of 64-bit IID will either specify explicitly how the "u" and "g" bits are set, or will simply include them as part of a field within the IID. In either case, a semantic meaning for these bits MUST NOT be defined.

In the following statement in section 2.5.1 of the IPv6 addressing architecture [RFC4291], the reference to "Modified EUI-64 format" applies only to cases where there is an underlying IEEE address:

"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."

The following statement in section 2.5.1 of the IPv6 addressing architecture [RFC4291] is hereby 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.

#### 5. Security Considerations

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

## 6. 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]. 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.

## 7. Acknowledgements

Valuable comments were received from Remi Despres, Fernando Gont, Brian Haberman, Joel Halpern, 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].

## 8. Change log [RFC Editor: Please remove]

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.

## 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.
- [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.

## 9.2. Informative References

- [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-03 (work in progress), January 2013.
- [I-D.ietf-softwire-4rd]  
Jiang, S., Despres, R., Penno, R., Lee, Y., Chen, G., and M. Chen, "IPv4 Residual Deployment via IPv6 - a Stateless Solution (4rd)", draft-ietf-softwire-4rd-04 (work in progress), October 2012.
- [RFC2526] Johnson, D. and S. Deering, "Reserved IPv6 Subnet Anycast Addresses", RFC 2526, March 1999.
- [RFC2629] Rose, M., "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.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.
- [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.
- [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](mailto: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](mailto:jiangsheng@huawei.com)

