Representing IPv6 Zone Identifiers in Uniform Resource Identifiers
draft-ietf-6man-uri-zoneid-00

Abstract

   This document describes how the Zone Identifier of an IPv6 scoped
   address can be represented in a Uniform Resource Identifier that
   includes a literal IPv6 address. It updates RFC 3986 and RFC 4007.

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 20, 2012.

Copyright Notice

Table of Contents

1. Introduction

   [RFC3986] defined how a literal IPv6 address can be represented in
   the "host" part of a Uniform Resource Identifier (URI).
   Subsequently, [RFC4007] extended the text representation of limited-
   scope IPv6 addresses such that a zone identifier may be concatenated
   to an address, for purposes described in that RFC. Zone identifiers
   are especially useful in contexts where literal addresses are
   typically used, for example during fault diagnosis, when it may be
   essential to specify which interface is used for sending to a link
   local address. It should be noted that zone identifiers have purely
   local meaning within the host where they are defined, and they are
   completely meaningless for any other host.

   RFC 4007 does not specify how zone identifiers are to be represented
   in URIs. Practical experience has shown that this feature is useful,
   in particular when using a web browser for debugging with link local
   addresses, but as it is undefined, it is not implemented consistently
   in URI parsers or in browsers.

   This document updates [RFC3986] by adding syntax to allow a zone
   identifier to be included in a literal IPv6 address. It also
   clarifies some statements in [RFC4007].

   It should be noted that in other contexts than a user interface, a
   zone identifier is mapped into a numeric zone index or interface
   number. The MIB textual convention [RFC4001] and the socket
   interface [RFC3493] define this as a 32 bit unsigned integer. The
   mapping between the human-readable zone identifier string and the
   numeric value is a host-specific function that varies between
   operating systems. The present document is concerned only with the
   human-readable string.

   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. Specification

   According to RFC 4007, a zone identifier is attached to the textual
   representation of an IPv6 address by concatenating "%" followed by
   <zone_id>, where <zone_id> is a string identifying the zone of the
   address. However, RFC 4007 gives no precise definition of the
   character set allowed in <zone_id>. There are no rules or de facto
   standards for this. For example, the first Ethernet interface in a
   host might be called %0, %1, %en1, %eth0, or whatever the implementer
   happened to choose.

In a URI, a literal IPv6 address is always embedded between "[" and
"]". This document specifies that <zone_id> may contain any ASCII
character classified in RFC 3986 as "unreserved", which conveniently
excludes "]" in order to simplify parsing.

There is an additional complication in that "%" is always treated as
an escape character in a URI, and according to RFC 3986 it MUST
therefore itself be escaped in a URI, in the form "%25". Thus, the
scoped address fe80::a%en1 would appear in a produced URI as
http://[fe80::a%25en1].

If an operating system uses any other characters in zone or interface
identifiers that are not in the "unreserved" character set, they too
MUST be escaped with a "%" sign according to RFC 3986.

In RFC 3986, the IPv6 literal format is formally defined in ABNF
[RFC5234] by the following rule:

    IP-literal = "[" (IPv6address / IPvFuture) "]"

To provide support for a zone identifier, the existing syntax of
IPv6address is retained, and a zone identifier may be added
optionally to any literal address. This allows flexibility for
unknown future uses. The rule quoted above from RFC 3986 is replaced
by three rules:

    IP-literal = "[" (IPv6addrz / IPvFuture) "]"

    ZoneID = 1*(unreserved / pct-encoded)

    IPv6addrz = IPv6address ["%" ZoneID]

The rules in [RFC5952] SHOULD be applied in producing URIs. The user
MUST replace "%" by "%25" when manually constructing such a URI, and
similarly for any other characters that are not in the "unreserved"
character set of RFC 3986. The option to include such characters
SHOULD NOT be used, but is included for the case of operating systems
that allow such characters in interface names.

The 6man WG discussed and rejected an alternative in which the
existing syntax of IPv6address would be extended by an option to add
the ZoneID only for the case of link-local addresses. It was felt
that the present solution offers more flexibility for future uses and
is more straightforward to implement.

RFC 4007 offers guidance on how the ZoneID affects interface/address
selection inside the IPv6 stack. Note that the behaviour of an IPv6
stack if passed a non-zero zone index for an address other than link-

   local is undefined.

3. Security Considerations

   The security considerations of [RFC3986] and [RFC4007] apply. In
   particular, this URI format creates a specific pathway by which a
   deceitful zone index might be communicated, as mentioned in the final
   security consideration of RFC 4007. It is emphasised that the format
   is intended only for debugging purposes, but of course this intention
   does not prevent misuse.

   To limit this risk, implementations SHOULD NOT allow use of this
   format except for well-defined usages such as sending to link local
   addresses under prefix fe80::/10.

   An HTTP server or proxy MUST ignore any ZoneID attached to an
   incoming URI, as it only has local significance at the sending host.

4. IANA Considerations

   This document requests no action by IANA.

5. Acknowledgements

   The lack of this format was pointed out by Kerry Lynn. Valuable
   comments and contributions were made by Karl Auer, Brian Haberman,
   Tatuya Jinmei, Tom Petch, Tomoyuki Sahara, and Juergen Schoenwaelder.

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

6. Change log [RFC Editor: Please remove]

   draft-ietf-6man-uri-zoneid-00: adopted by WG, fixed syntax to allow
   for % encoded characters, 2012-02-17.

   draft-carpenter-v6ops-label-balance-01: chose Option 2, removed 15
   character limit, added explanation of ID/number mapping and other
   clarifications, 2012-02-08.

   draft-carpenter-v6ops-label-balance-00: original version, 2011-12-07.

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.

[RFC3986]   Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform
            Resource Identifier (URI): Generic Syntax", STD 66,
            RFC 3986, January 2005.

[RFC4007]   Deering, S., Haberman, B., Jinmei, T., Nordmark, E., and
            B. Zill, "IPv6 Scoped Address Architecture", RFC 4007,
            March 2005.

[RFC5234]   Crocker, D. and P. Overell, "Augmented BNF for Syntax
            Specifications: ABNF", STD 68, RFC 5234, January 2008.

[RFC5952]   Kawamura, S. and M. Kawashima, "A Recommendation for IPv6
            Address Text Representation", RFC 5952, August 2010.

[7.2](). Informative References

[RFC2629]   Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629,
            June 1999.

[RFC3493]   Gilligan, R., Thomson, S., Bound, J., McCann, J., and W.
            Stevens, "Basic Socket Interface Extensions for IPv6",
            RFC 3493, February 2003.

[RFC4001]   Daniele, M., Haberman, B., Routhier, S., and J.
            Schoenwaelder, "Textual Conventions for Internet Network
            Addresses", RFC 4001, February 2005.

Authors' Addresses

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

    Email: brian.e.carpenter@gmail.com

      Robert M. Hinden
      Check Point Software Technologies, Inc.
      800 Bridge Parkway
      Redwood City, CA 94065
      US

      Email: bob.hinden@gmail.com