

An ".ipv6" Top Level Pseudo-Domain

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC 2026](#).

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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.''

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

Abstract

The normal textual representation for IPv6 addresses as a set of colon-separated hexadecimal numbers does not work well with most deployed URL-parsing software. This document describes an alternate format which will pass unharmed through most URL-parsing software.

1. Introduction

The normal textual representation for IPv6 addresses as a set of colon-separated hexadecimal numbers does not work well with most deployed URL-parsing software. This document describes an alternate format which will pass unharmed through most URL-parsing software.

Rather than describing an alternative syntax for IPv6 addresses only within URLs, it proposes a pseudo-domain ".ipv6" for using a domain-name-like syntax that actually represents the literal IPv6 address.

2. Background

The standard representation for IPv6 addresses in text is defined in [section 2.2 of \[RFC2373\]](#) ("Text Representation of Addresses"). This representation uses hexadecimal values separated by colon ":", double colon "::", and optionally ending period-separated decimal values for the four low-order 8-bit pieces of the address.

Unfortunately, using this IPv6 syntax within URLs [\[RFC2396\]](#) would be disruptive for many applications. Within the "hostport" section of the generic URI syntax, the colon is used to separate the host name or address from an (optional) port number. Thus, in some addresses, a colon followed by a decimal number could ambiguously be interpreted as a port designator or as a part of the IPv6 address.

Even if there were no ambiguity, this syntax is incompatible with a many deployed applications that parse (but do not resolve) URLs, including many CGI scripts, robots, search engines, and so forth.

3. Syntax

This specification defines a simple, safe representation for IPv6 addresses which can be used within URLs and also in other contexts. It does so by defining a syntax which will look like a domain name to otherwise unaware software.

The syntax is best described as a transformation of the normal IPv6 syntax as defined in [section 2.2 of \[RFC2373\]](#); starting with such an address:

- 1) replace every colon ":" with a "-"
- 2) append ".ipv6" to the end

Thus, an HTTP service available at port 70 of IPv6 address "ABCD:EF01::2345:10.9.8.7" could be written as

<http://ABCD-EF01--2345-10.9.8.7.ipv6:70/>

This syntax should always be used. Internet software that resolves host names and addresses in URLs should be modified to recognize the "ipv6" pseudo-domain.

4. Implementation considerations

It would be useful for systems to uniformly recognize the "ipv6" pseudo-domain throughout, e.g., as a special case within "gethostbyname". [\[RFC 2133\]](#).

5. IANA considerations

IANA is requested to reserve the "ipv6" top level pseudo-domain for the purpose outlined in this memo.

6. References

- [RFC2396] R. Fielding, L. Masinter, T. Berners-Lee, "Uniform Resource Identifiers: Generic Syntax", August, 1998.
- [RFC2373] R. Hinden, S. Deering. "IP Version 6 Addressing Architecture", July, 1998.
- [RFC2133] R. Gilligan, S. Thomson, J. Bound, W.Stevens.
"Basic Socket Interface Extensions for IPv6".

7. Authors' Addresses

Larry Masinter
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94034, USA
Fax: +1 650 812 4365
EMail: masinter@parc.xerox.com

James Gettys
MIT Laboratory for Computer Science
545 Technology Square
Cambridge, MA 02139, USA
Fax: +1 617 258 8682
Email: jg@w3.org

Brian Carpenter
IBM United Kingdom Laboratories
MP 185, Hursley Park
Winchester, Hampshire SO21 2JN, UK
Email: brian@hursley.ibm.com

7. Full Copyright Statement

Copyright (C) The Internet Society (1999). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process

must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.