

Internet Engineering Task Force (IETF)
Request for Comments: 5954
Updates: 3261
Category: Standards Track
ISSN: 2070-1721

V. Gurbani, Ed.
Bell Laboratories, Alcatel-Lucent
B. Carpenter, Ed.
Univ. of Auckland
B. Tate, Ed.
BroadSoft
August 2010

Essential Correction for IPv6 ABNF and URI Comparison in RFC 3261

Abstract

This document corrects the Augmented Backus-Naur Form (ABNF) production rule associated with generating IPv6 literals in RFC 3261. It also clarifies the rule for Uniform Resource Identifier (URI) comparison when the URIs contain textual representation of IP addresses.

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/rfc5954>.

Copyright Notice

Copyright (c) 2010 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
- 2. Terminology 2
- 3. Problem Statement 2
 - 3.1. Extra Colon in IPv4-Mapped IPv6 Address 2
 - 3.2. Comparing URIs with Textual Representation of IP Addresses 3
- 4. Resolution 4
 - 4.1. Resolution for Extra Colon in IPv4-Mapped IPv6 Address . . 4
 - 4.2. Clarification for Comparison of URIs with Textual Representation of IP Addresses 5
- 5. Generating a Canonical IPv6 Textual Representation 5
- 6. Security Considerations 6
- 7. Acknowledgments 6
- 8. References 6
 - 8.1. Normative References 6
 - 8.2. Informative References 6

1. Introduction

This document corrects the Augmented Backus-Naur Form (ABNF) production rule associated with generating IPv6 literals in RFC 3261 [1]. It also clarifies the rule for Uniform Resource Identifier (URI) comparison when the URIs contain textual representation of IP addresses.

2. 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 RFC 2119 [2].

3. Problem Statement

3.1. Extra Colon in IPv4-Mapped IPv6 Address

The ABNF [4] for generating IPv6 literals in RFC 3261 [1] is incorrect. When generating IPv4-mapped IPv6 addresses, the production rule may actually generate the following construct:

[2001:db8:::192.0.2.1] - Note the extra colon before the IPv4 address.

The correct construct, of course, would only include two colons before the IPv4 address.

Historically, the ABNF pertaining to IPv6 references in RFC 3261 was derived from Appendix B of RFC 2373 [7], which was flawed to begin with (see errata for RFC 2373 [8]). RFC 2373 has been subsequently obsoleted by RFC 4291 [6].

The ABNF for IPv6reference is reproduced from RFC 3261 below:

```

IPv6reference = "[" IPv6address "]"
IPv6address  = hexpart [ ":" IPv4address ]
IPv4address  = 1*3DIGIT "." 1*3DIGIT "." 1*3DIGIT "." 1*3DIGIT
hexpart      = hexseq / hexseq ":" [ hexseq ] / ":" [ hexseq ]
hexseq       = hex4 *( ":" hex4)
hex4         = 1*4HEXDIG

```

Note that the ambiguity occurs in the <IPv6address> production rule where the <IPv4address> non-terminal is prefixed by the ":" token. Because the <hexpart> production rule is defined such that two of its alternatives already include the ":" token, this may yield to the faulty construction of an IPv6-mapped IPv4 address with an extra colon when expanding those alternatives.

3.2. Comparing URIs with Textual Representation of IP Addresses

In SIP, URIs are compared for a variety of reasons. Registrars compare URIs when they receive a binding update request, for instance. Section 19.1.4 of RFC 3261 [1] provides the rules for comparing URIs. Among other rules, it states that:

For two URIs to be equal, the user, password, host, and port components must match.

Does the above rule then imply that the following URIs are equal:

`sip:bob@[::ffff:192.0.2.128] = sip:bob@[::ffff:c000:280]?`

`sip:bob@[2001:db8::9:1] = sip:bob@[2001:db8::9:01]?`

`sip:bob@[0:0:0:0:0:FFFF:129.144.52.38] = sip:bob@[::FFFF:129.144.52.38]?`

In all of the above examples, the textual representation of the IPv6 address is different, but these addresses are binary equivalents (implementers are also urged to consult Section 5 of this document for recommendations on IPv6 address text representations). Section 19.1.4 of RFC 3261 does not provide any rule for URIs containing different textual representations of IPv6 addresses that all correspond to the same binary equivalent.

Note that the same ambiguity occurs for IPv4 addresses, i.e., is 192.0.2.128 = 192.00.02.128? However, IPv6, with its compressed notation and the need to represent hybrid addresses (like IPv4-mapped IPv6 addresses) makes the representation issue more acute. The resolution discussed in Section 4.2 applies to textual representations of both IPv6 and IPv4 addresses.

4. Resolution

4.1. Resolution for Extra Colon in IPv4-Mapped IPv6 Address

The resolution to this ambiguity is simply to use the correct ABNF for the <IPv6address> production rule from Appendix A of RFC 3986 [3]. For the sake of completeness, it is reproduced below:

```

IPv6address      =
  /
  / [
  / [ *1( h16 ":" ) h16 ] ":" 4( h16 ":" ) ls32
  / [ *2( h16 ":" ) h16 ] ":" 3( h16 ":" ) ls32
  / [ *3( h16 ":" ) h16 ] ":" 2( h16 ":" ) ls32
  / [ *4( h16 ":" ) h16 ] ":"   h16 ":"   ls32
  / [ *5( h16 ":" ) h16 ] ":"           ls32
  / [ *6( h16 ":" ) h16 ] ":"           h16

h16              = 1*4HEXDIG
ls32             = ( h16 ":" h16 ) / IPv4address
IPv4address      = dec-octet "." dec-octet "." dec-octet "." dec-octet
dec-octet        = DIGIT ; 0-9
                  / %x31-39 DIGIT ; 10-99
                  / "1" 2DIGIT ; 100-199
                  / "2" %x30-34 DIGIT ; 200-249
                  / "25" %x30-35 ; 250-255

```

Accordingly, this document updates RFC 3261 as follows: the <IPv6address> and <IPv4address> production rules from RFC 3261 MUST NOT be used and instead, the production rules of the same name in RFC 3986 (and reproduced above) MUST be used. This will render <hexpart>, <hexseq>, and <hex4> production rules in RFC 3261 obsolete; as such, these three production rules -- namely, <hexpart>, <hexseq>, and <hex4> -- from RFC 3261 MUST NOT be used.

The use of the <IPv4address> production rule from RFC 3986 no longer allows syntactically valid -- though semantically invalid -- SIP URIs of the form "sip:bob@444.555.666.777".

4.2. Clarification for Comparison of URIs with Textual Representation of IP Addresses

The resolution to this ambiguity is a simple clarification acknowledging that the textual representation of an IP address varies, but it is the binary equivalence of the IP address that must be taken into consideration when comparing two URIs that contain varying textual representations of an IP address.

Accordingly, the existing rule from the bulleted list in Section 19.1.4 of RFC 3261 MUST be modified as follows:

OLD:

- o For two URIs to be equal, the user, password, host, and port components must match.

NEW:

- o For two URIs to be equal, the user, password, host, and port components must match. If the host component contains a textual representation of IP addresses, then the representation of those IP addresses may vary. If so, the host components are considered to match if the different textual representations yield the same binary IP address.

In addition, the text in the following paragraph MUST be added to the existing list of examples in Section 19.1.4 of RFC 3261 in order to demonstrate the intent of the modified rule:

The following URIs are equivalent because the underlying binary representation of the IP addresses are the same although their textual representations vary:

```
sip:bob@[::ffff:192.0.2.128]
sip:bob@[::ffff:c000:280]
```

```
sip:bob@[2001:db8::9:1]
sip:bob@[2001:db8::9:01]
```

```
sip:bob@[0:0:0:0:0:FFFF:129.144.52.38]
sip:bob@[::FFFF:129.144.52.38]
```

5. Generating a Canonical IPv6 Textual Representation

Implementers SHOULD generate IPv6 text representation as defined in RFC 5952 [5].

6. Security Considerations

This document does not introduce any new security considerations beyond those described in RFC 3261 [1].

7. Acknowledgments

The ABNF for IPv6 was developed by Roy T. Fielding and Andrew Main and published in RFC 3986.

Jeroen van Bommel, Peter Blatherwick, Gonzalo Camarillo, Paul Kyzivat, Jonathan Rosenberg, Michael Thomas, and Dale Worley provided invaluable discussion points on the SIP WG mailing list on the URI equivalency problem. Alfred Hoenes urged the use of angle brackets (as specified in Section 2.1 of RFC 5234 [4]) to denote productions.

8. References

8.1. Normative References

- [1] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [3] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [4] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [5] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, August 2010.

8.2. Informative References

- [6] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [7] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, July 1998.
- [8] "RFC Editor Errata", <<http://www.rfc-editor.org/errata.php>>.

Authors' Addresses

Vijay K. Gurbani (editor)
Bell Laboratories, Alcatel-Lucent
1960 Lucent Lane
Room 9C-533
Naperville, IL 60563
USA

Phone: +1 630 224-0216
EMail: vkg@bell-labs.com

Brian E. Carpenter (editor)
Department of Computer Science
University of Auckland
PB 92019
Auckland, 1142
New Zealand

EMail: brian.e.carpenter@gmail.com

Brett Tate (editor)
BroadSoft

EMail: brett@broadsoft.com