

MIB Textual Conventions for Uniform Resource Identifiers (URIs)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This MIB module defines textual conventions to represent STD 66 Uniform Resource Identifiers (URIs). The intent is that these textual conventions will be imported and used in MIB modules that would otherwise define their own representation(s).

Table of Contents

1. Introduction	1
2. Terminology	2
3. The Internet-Standard Management Framework	2
4. Definitions	2
5. Security Considerations	5
6. IANA Considerations	5
7. Acknowledgements	5
8. References	5
8.1. Normative References	5
8.2. Informative References	6

1. Introduction

This memo defines a portion of the Management Information Base (MIB) for use with network management protocols in the Internet community. It defines textual conventions to represent STD 66 [RFC3986] URIs, which are further described by [RFC3305].

Three textual conventions are defined: one of unrestricted length, and two of different restricted lengths. Which length is appropriate will depend on tradeoffs made in particular MIB modules. The purpose of providing standard restricted-length textual conventions is to improve compatibility between MIB modules that require restricted-length URIs.

If a URI needs to be used as an index object, then the 'Uri' TEXTUAL-CONVENTION SHOULD be subtyped to a length appropriate for the Object Identifier (OID) of which it is part. The description of the 'Uri' TEXTUAL-CONVENTION discusses this case.

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 [RFC2119].

3. The Internet-Standard Management Framework

For a detailed overview of the documents that describe the current Internet-Standard Management Framework, please refer to section 7 of RFC 3410 [RFC3410].

Managed objects are accessed via a virtual information store, termed the Management Information Base or MIB. MIB objects are generally accessed through the Simple Network Management Protocol (SNMP). Objects in the MIB are defined using the mechanisms defined in the Structure of Management Information (SMI). This memo specifies a MIB module that is compliant to the SMIV2, which is described in STD 58, RFC 2578 [RFC2578], STD 58, RFC 2579 [RFC2579] and STD 58, RFC 2580 [RFC2580].

4. Definitions

URI-TC-MIB DEFINITIONS ::= BEGIN

IMPORTS

MODULE-IDENTITY, mib-2	FROM SNMPv2-SMI	-- [RFC2578]
TEXTUAL-CONVENTION	FROM SNMPv2-TC;	-- [RFC2579]

uriTcMIB MODULE-IDENTITY

LAST-UPDATED "200709100000Z" -- 10 September 2007
ORGANIZATION "IETF Operations and Management (OPS) Area"
CONTACT-INFO "EMail: ops-area@ietf.org
 Home page: <http://www.ops.ietf.org/>"
DESCRIPTION
 "This MIB module defines textual conventions for
 representing URIs, as defined by RFC 3986 STD 66."
REVISION "200709100000Z" -- 10 September 2007
DESCRIPTION
 "Initial revision, published as RFC 5017."

Copyright (C) The IETF Trust (2007). This version of this
MIB module is part of RFC 5017; see the RFC itself for full

```
        legal notices."  
 ::= { mib-2 164 }
```

Uri ::= TEXTUAL-CONVENTION

DISPLAY-HINT "1a"

STATUS current

DESCRIPTION

"A Uniform Resource Identifier (URI) as defined by STD 66.

Objects using this TEXTUAL-CONVENTION MUST be in US-ASCII encoding, and MUST be normalized as described by RFC 3986 Sections 6.2.1, 6.2.2.1, and 6.2.2.2. All unnecessary percent-encoding is removed, and all case-insensitive characters are set to lowercase except for hexadecimal digits, which are normalized to uppercase as described in Section 6.2.2.1.

The purpose of this normalization is to help provide unique URIs. Note that this normalization is not sufficient to provide uniqueness. Two URIs that are textually distinct after this normalization may still be equivalent.

Objects using this TEXTUAL-CONVENTION MAY restrict the schemes that they permit. For example, 'data:' and 'urn:' schemes might not be appropriate.

A zero-length URI is not a valid URI. This can be used to express 'URI absent' where required, for example when used as an index field.

Where this TEXTUAL-CONVENTION is used for an index field, it MUST be subtyped to restrict its length. There is an absolute limit of 128 subids for an OID, and it is not efficient to have OIDs whose length approaches this limit."

REFERENCE "RFC 3986 STD 66 and RFC 3305"

SYNTAX OCTET STRING

Uri255 ::= TEXTUAL-CONVENTION

DISPLAY-HINT "255a"

STATUS current

DESCRIPTION

"A Uniform Resource Identifier (URI) as defined by STD 66.

Objects using this TEXTUAL-CONVENTION MUST be in US-ASCII encoding, and MUST be normalized as described by RFC 3986 Sections 6.2.1, 6.2.2.1, and 6.2.2.2. All unnecessary percent-encoding is removed, and all case-insensitive

characters are set to lowercase except for hexadecimal digits, which are normalized to uppercase as described in Section 6.2.2.1.

The purpose of this normalization is to help provide unique URIs. Note that this normalization is not sufficient to provide uniqueness. Two URIs that are textually distinct after this normalization may still be equivalent.

Objects using this TEXTUAL-CONVENTION MAY restrict the schemes that they permit. For example, 'data:' and 'urn:' schemes might not be appropriate.

A zero-length URI is not a valid URI. This can be used to express 'URI absent' where required, for example when used as an index field.

STD 66 URIs are of unlimited length. Objects using this TEXTUAL-CONVENTION impose a length limit on the URIs that they can represent. Where no length restriction is required, objects SHOULD use the 'Uri' TEXTUAL-CONVENTION instead. Objects used as indices SHOULD subtype the 'Uri' TEXTUAL-CONVENTION."

REFERENCE "RFC 3986 STD 66 and RFC 3305"
SYNTAX OCTET STRING (SIZE (0..255))

Uri1024 ::= TEXTUAL-CONVENTION

DISPLAY-HINT "1024a"

STATUS current

DESCRIPTION

"A Uniform Resource Identifier (URI) as defined by STD 66.

Objects using this TEXTUAL-CONVENTION MUST be in US-ASCII encoding, and MUST be normalized as described by RFC 3986 Sections 6.2.1, 6.2.2.1, and 6.2.2.2. All unnecessary percent-encoding is removed, and all case-insensitive characters are set to lowercase except for hexadecimal digits, which are normalized to uppercase as described in Section 6.2.2.1.

The purpose of this normalization is to help provide unique URIs. Note that this normalization is not sufficient to provide uniqueness. Two URIs that are textually distinct after this normalization may still be equivalent.

Objects using this TEXTUAL-CONVENTION MAY restrict the schemes that they permit. For example, 'data:' and 'urn:' schemes might not be appropriate.

A zero-length URI is not a valid URI. This can be used to express 'URI absent' where required, for example when used as an index field.

STD 66 URIs are of unlimited length. Objects using this TEXTUAL-CONVENTION impose a length limit on the URIs that they can represent. Where no length restriction is required, objects SHOULD use the 'Uri' TEXTUAL-CONVENTION instead. Objects used as indices SHOULD subtype the 'Uri' TEXTUAL-CONVENTION."

REFERENCE "RFC 3986 STD 66 and RFC 3305"

SYNTAX OCTET STRING (SIZE (0..1024))

END

5. Security Considerations

See also the Security Considerations of STD 66 [RFC3986].

This MIB module does not define any management objects. Instead, it defines a textual convention that may be imported by other MIB modules and used for object definitions.

Meaningful security considerations can only be written in the MIB modules that define management objects. This document therefore has no impact on the security of the Internet.

6. IANA Considerations

URI-TC-MIB is rooted under the mib-2 subtree. IANA has assigned { mib-2 164 } to the URI-TC-MIB module specified in this document.

7. Acknowledgements

This module was generated by editing together contributions from Randy Presuhn, Dan Romascanu, Bill Fenner, Juergen Schoenwaelder, and others.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2578] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Structure of Management Information Version 2 (SMIv2)", STD 58, RFC 2578, April 1999.

- [RFC2579] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Textual Conventions for SMIV2", STD 58, RFC 2579, April 1999.
- [RFC2580] McCloghrie, K., Perkins, D., and J. Schoenwaelder, "Conformance Statements for SMIV2", STD 58, RFC 2580, April 1999.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.

8.2. Informative References

- [RFC3305] Mealling, M. and R. Denenberg, "Report from the Joint W3C/IETF URI Planning Interest Group: Uniform Resource Identifiers (URIs), URLs, and Uniform Resource Names (URNs): Clarifications and Recommendations", RFC 3305, August 2002.
- [RFC3410] Case, J., Mundy, R., Partain, D., and B. Stewart, "Introduction and Applicability Statements for Internet-Standard Management Framework", RFC 3410, December 2002.

Author's Address

David McWalter (editor)
Data Connection Ltd
100 Church Street
Enfield EN2 6BQ
United Kingdom

EMail: dmcw@dataconnection.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

