

Internet Engineering Task Force (IETF)
Request for Comments: 5935
Category: Standards Track
ISSN: 2070-1721

M. Ellison
Ellison Software Consulting
B. Natale
MITRE
August 2010

Expressing SNMP SMI Datatypes in XML Schema Definition Language

Abstract

This memo defines the IETF standard expression of Structure of Management Information (SMI) base datatypes in XML Schema Definition (XSD) language. The primary objective of this memo is to enable the production of XML documents that are as faithful to the SMI as possible, using XSD as the validation mechanism.

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/rfc5935>.

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. Conventions	4
3. Requirements	4
4. XSD for SMI Base Datatypes	5
5. Rationale	8
5.1. Numeric Datatypes	8
5.2. OctetString	9
5.3. Opaque	10
5.4. IPAddress	10
5.5. ObjectIdentifier	10
6. Security Considerations	11
7. IANA Considerations	11
7.1. SMI Base Datatypes Namespace Registration	12
7.2. SMI Base Datatypes Schema Registration	12
8. Acknowledgements	12
9. References	13
9.1. Normative References	13
9.2. Informative References	13

1. Introduction

Numerous use cases exist for expressing the management information described by SMI Management Information Base (MIB) modules in XML [XML]. Potential use cases reside both outside and within the traditional IETF network management community. For example, developers of some XML-based management applications may want to incorporate the rich set of data models provided by MIB modules. Developers of other XML-based management applications may want to access MIB module instrumentation via gateways to SNMP agents. Such applications benefit from the IETF standard mapping of SMI datatypes to XML datatypes via XSD [XMLSchema], [XSDDatatypes].

MIB modules use SMIV2 [RFC2578] to describe data models. For legacy MIB modules, SMIV1 [RFC1155] was used. MIB data conveyed in variable bindings ("varbinds") within protocol data units (PDUs) of SNMP messages use the primitive, base datatypes defined by the SMI.

The SMI allows for the creation of derivative datatypes, "textual conventions" ("TCs") [RFC2579]. A TC has a unique name, has a syntax that either refines or is a base SMI datatype, and has relatively precise application-level semantics. TCs facilitate correct application-level handling of MIB data, improve readability of MIB modules by humans, and support appropriate renderings of MIB data.

Values in varbinds corresponding to MIB objects defined with TC syntax are always encoded as the base SMI datatype underlying the TC syntax. Thus, the XSD mappings defined in this memo provide support for values of MIB objects defined with TC syntax as well as for values of MIB objects defined with base SMI syntax. Using the translation of TC into base SMI datatypes any MIB module that uses TCs can be mapped into XSD using the mappings defined in this memo. For example, for IP addresses (both IPv4 and IPv6), MIB objects defined using the InetAddress TC (as per [RFC4001]) are encoded using the base SMI datatype underlying the InetAddress TC syntax rather than the IPAddress base datatype.

Various independent schemes have been devised for expressing SMI datatypes in XSD. These schemes exhibit a degree of commonality, especially concerning numeric SMI datatypes, but these schemes also exhibit sufficient differences, especially concerning the non-numeric SMI datatypes, precluding uniformity of expression and general interoperability.

Throughout this memo, the term "fidelity" refers to the quality of an accurate, consistent representation of SMI data values and the term "faithful" refers to the quality of reliably reflecting the semantics of SMI data values. Thus defined, the characteristics of fidelity and being faithful are essential to uniformity of expression and general interoperability in the XML representation of SMI data values.

The primary purpose of this memo is to define the standard expression of SMI base datatypes in XML documents that is both uniform and interoperable. This standard expression enables Internet operators, management application developers, and users to benefit from a wider range of management tools and to benefit from a greater degree of unified management. Thus, standard expression enables and facilitates improvements to the timeliness, accuracy, and utility of management information.

The overall objective of this memo, and of any related future memos as may be published, is to define the XSD equivalent [XSDDatatypes] of SMIV2 (STD 58) and to encourage XML-based protocols to carry, and XML-based applications to use, the management information defined in SMIV2-compliant MIB modules. The use of a standard mapping from SMIV2 to XML via XSD validation enables and promotes the efficient reuse of existing and future MIB modules and instrumentation by XML-based protocols and management applications.

Developers of certain XML-based management applications will find this specification sufficient for their purposes. Developers of other XML-based management applications may need to make more

complete reuse of existing MIB modules, requiring standard XSD documents for TCs [RFC2579] and MIB structure [RFC2578]. Memos supporting such requirements are planned, but have not been produced at the time of this writing.

Finally, it is worthwhile to note that the goal of fidelity to the SMIV2 standard (STD 58), as specified in the "Requirements" section below, is crucial to this effort. Fidelity leverages the established "rough consensus" of the precise SMIV2 data models contained in MIB modules, and leverages existing instrumentation, the "running code" implementing SMIV2 data models. This effort does not include any redesign of SMIV2 datatypes, data structures or textual conventions in order to overcome known limitations. Such work can be pursued by other efforts.

2. Conventions

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. Requirements

The following set of requirements is intended to produce XML documents that can be validated via the XSD defined in this specification to faithfully represent values carried "on-the-wire" in SNMP PDUs as defined by the SMI:

- R1. All SMI base datatypes MUST have a corresponding XSD datatype.
- R2. SMIV2 is the normative SMI for this document. Prior to mapping datatypes into XSD, legacy SMIV1 modules MUST be converted (at least logically) in accordance with Section 2.1, inclusive, of the "Coexistence" RFC [RFC3584].
- R3. The XSD datatype specified for a given SMI datatype MUST be able to represent all valid values for that SMI datatype.
- R4. The XSD datatype specified for a given SMI datatype MUST represent any special encoding rules associated with that SMI datatype.
- R5. The XSD datatype specified for a given SMI datatype MUST include any restrictions on values associated with the SMI datatype.

- R6. The XSD datatype specified for a given SMI datatype MUST be the most logical XSD datatype, with the fewest necessary restrictions on its set of values, consistent with the foregoing requirements.
- R7. The XML output produced as a result of meeting the foregoing requirements SHOULD be the most coherent and succinct representation (i.e., avoiding superfluous "decoration") from the perspective of readability by humans.
4. XSD for SMI Base Datatypes

This document provides XSD datatype mappings for the SMIV2 base datatypes only -- i.e., the eleven "ObjectSyntax" datatypes defined in RFC 2578. These datatypes -- via tag values defined in the SMIV2 to identify them in varbinds -- constrain values carried "on-the-wire" in SNMP PDUs between SNMP management applications and SNMP agents:

- o INTEGER, Integer32
- o Unsigned32, Gauge32
- o Counter32
- o TimeTicks
- o Counter64
- o OCTET STRING
- o Opaque
- o IpAddress
- o OBJECT IDENTIFIER

The "BITS" pseudo-type (also referred to as a "construct" in RFC 2578) is treated as a Textual Convention, not a base datatype, for the purpose of this document.

BEGIN

```
<?xml version="1.0" encoding="utf-8"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
xmlns="urn:ietf:params:xml:ns:smi:base:1.0"
targetNamespace="urn:ietf:params:xml:ns:smi:base:1.0"
elementFormDefault="qualified"
attributeFormDefault="unqualified"
xml:lang="en">
```

```
<xs:annotation>
```

```
<xs:documentation>
```

```
Mapping of SMIV2 base datatypes from RFC 2578
```

```
Contact:      Mark Ellison
Organization: Ellison Software Consulting
Address:      38 Salem Road
              Atkinson, NH 03811
              USA
Telephone:    +1 603-362-9270
E-Mail:       ietf@EllisonSoftware.com
```

```
Contact:      Bob Natale
Organization: MITRE
Address:      300 Sentinel Drive
              6th Floor
              Annapolis Junction, MD 20701
              USA
Telephone:    +1 301-617-3008
E-Mail:       rnatale@mitre.org
```

```
Last Updated: 201002260000Z
```

```
Copyright (c) 2010 IETF Trust and the persons identified as
authors of the code. All rights reserved.
```

```
Redistribution and use in source and binary forms, with or
without modification, is permitted pursuant to, and subject
to the license terms contained in, the Simplified BSD License
set forth in Section 4.c of the IETF Trust's Legal Provisions
Relating to IETF Documents
(http://trustee.ietf.org/license-info).
```

```
This version of this XML Schema Definition (XSD)
document is part of RFC 5935; see the RFC itself for
full legal notices.
```

```
</xs:documentation>
```

```
</xs:annotation>

<xs:simpleType name="INTEGER">
  <xs:restriction base="xs:int"/>
</xs:simpleType>

<xs:simpleType name="Integer32">
  <xs:restriction base="xs:int"/>
</xs:simpleType>

<xs:simpleType name="Unsigned32">
  <xs:restriction base="xs:unsignedInt"/>
</xs:simpleType>

<xs:simpleType name="Gauge32">
  <xs:restriction base="xs:unsignedInt"/>
</xs:simpleType>

<xs:simpleType name="Counter32">
  <xs:restriction base="xs:unsignedInt"/>
</xs:simpleType>

<xs:simpleType name="TimeTicks">
  <xs:restriction base="xs:unsignedInt"/>
</xs:simpleType>

<xs:simpleType name="Counter64">
  <xs:restriction base="xs:unsignedLong"/>
</xs:simpleType>

<xs:simpleType name="OctetString">
  <xs:restriction base="xs:hexBinary">
    <xs:maxLength value="65535"/>
  </xs:restriction>
</xs:simpleType>

<xs:simpleType name="Opaque">
  <xs:restriction base="xs:hexBinary"/>
</xs:simpleType>

<xs:simpleType name="IpAddress">
  <xs:restriction base="xs:string">
    <xs:pattern value=
      "(([0-9]|[1-9][0-9]|1[0-9][0-9]|2[0-4][0-9]|25[0-5])\.){3}
      ([0-9]|[1-9][0-9]|1[0-9][0-9]|2[0-4][0-9]|25[0-5])"/>
  </xs:restriction>
</xs:simpleType>
```

```
<xs:simpleType name="ObjectIdentifier">
  <xs:restriction base="xs:string">
    <xs:pattern value=
      "([0-1](\.[1-3]?[0-9]))|
      (2\.(0|([1-9]\d*)))|
      (\.(0|([1-9]\d*))) {0,126}"/>
    </xs:restriction>
  </xs:simpleType>

</xs:schema>
END
```

5. Rationale

The XSD datatypes, including any specified restrictions, were chosen based on fit with the requirements specified earlier in this document, and with attention to simplicity while maintaining fidelity to the SMI. Also, the "canonical representations" (i.e., refinements of the "lexical representations") documented in the W3C XSD specification [XMLSchema], [XSDDatatypes] are assumed.

5.1. Numeric Datatypes

All of the numeric XSD datatypes specified in the previous section -- INTEGER, Integer32, Unsigned32, Gauge32, Counter32, TimeTicks, and Counter64 -- comply with the relevant requirements

- o They cover all valid values for the corresponding SMI datatypes.
- o They comply with the standard encoding rules associated with the corresponding SMI datatypes.
- o They inherently match the range restrictions associated with the corresponding SMI datatypes.
- o They are the most direct XSD datatypes that exhibit the foregoing characteristics relative to the corresponding SMI datatypes (which is why no "restriction" statements -- other than the "base" XSD type -- are required in the XSD).
- o The XML output produced from the canonical representation of these XSD datatypes is also the most direct from the perspective of readability by humans (i.e., no leading "+" sign and no leading zeros).

Special note to application developers: compliance with this schema in an otherwise correct translation from raw ("on-the-wire" representation) SNMP MIB data produces values that are faithful to the original. However, the Gauge32, Counter32, Counter64, and TimeTicks datatypes have special application semantics that must be considered when using their raw values for anything other than display, printing, storage, or transmission of the literal value. RFC 2578 provides the necessary details.

5.2. OctetString

This XSD datatype corresponds to the SMI "OCTET STRING" datatype.

Several independent schemes for mapping SMI datatypes to XSD have used the XSD "string" type to represent "OCTET STRING", but this mapping does not conform to the requirements specified in this document. Most notably, "string" cannot faithfully represent all valid values (0 thru 255) that each octet in an "OCTET STRING" can have -- or at least cannot do so in a way that provides for easy human readability of the resulting XML output.

Consequently, the XSD datatype "hexBinary" is specified as the standard mapping of the SMI "OCTET STRING" datatype. In hexBinary, each octet is encoded as two hexadecimal digits; the canonical representation limits the set of allowed hexadecimal digits to 0-9 and uppercase A-F.

The hexBinary representation of "OCTET STRING" complies with the relevant requirements:

- o It covers all valid values for the corresponding SMI datatype.
- o It complies with the standard encoding rules associated with the corresponding SMI datatype.
- o With the "maxLength" restriction to 65535 octets, the XSD datatype specification matches the restrictions associated with the corresponding SMI datatype.
- o It is the most direct XSD datatype that exhibits the foregoing characteristics relative to the corresponding SMI datatype (which must allow for any valid binary octet value).
- o The XML output produced from the canonical representation of this XSD datatype is not optimal with respect to readability by humans; however, that is a consequence of the SMI datatype itself. Where human readability is more of a concern, it is likely that the

actual MIB objects in question will be represented by textual conventions that limit the set of values that will be included in the OctetStrings and will, thus, bypass the hexBinary typing.

5.3. Opaque

The "hexBinary" XSD datatype is specified as the representation of the SMI "Opaque" datatype generally for the same reasons as "hexBinary" is specified for the "OctetString" datatype:

- o It covers all valid values for the corresponding SMI datatype.
- o It complies with the standard encoding rules associated with the corresponding SMI datatype.
- o There are no restriction issues associated with using "hexBinary" for "Opaque".
- o It is the most direct XSD datatype that exhibits the foregoing characteristics relative to the corresponding SMI datatype (which must allow for any valid binary octet value).
- o The XML output produced from the canonical representation of this XSD datatype is not optimal with respect to readability by humans; however, that is a consequence of the SMI datatype itself. Unmediated "Opaque" data is intended for consumption by applications, not humans.

5.4. IPAddress

The XSD "string" datatype is the natural choice to represent an IPAddress as XML output. The "pattern" restriction applied in this case results in a dotted-decimal string of four values between "0" and "255" separated by a period (".") character. This pattern also precludes leading zeros.

Note that the SMI relies upon Textual Conventions (TCs) to specify an IPv6 address. As such, the representation of an IPv6 address as an XSD datatype is beyond the scope of this document.

5.5. ObjectIdentifier

This XSD datatype corresponds to the SMI "OBJECT IDENTIFIER" datatype.

The XSD "string" datatype is also the natural choice to represent an ObjectIdentifier as XML output, for the same reasons as for the IPAddress choice. The "pattern" restriction applied in this case results in a dotted-decimal string of up to 128 elements (referred to as "sub-ids"), each holding an "Unsigned32" integer value.

Note that the first two components of an "OBJECT IDENTIFIER" each have a limited range of values as indicated in the XSD pattern restriction and as described in the ASN1.1/BER standard [ASN.1].

There are three values allocated for the root node, and at most 39 values for nodes subordinate to a root node value of 0 or 1.

The minimum length of an "OBJECT IDENTIFIER" is two sub-ids and the representation of a zero-valued "OBJECT IDENTIFIER" is "0.0".

Note that no explicit "minLength" restriction, which would be "3" to allow for the minimum of two sub-ids and a single separating dot, is required since the pattern itself enforces this restriction.

6. Security Considerations

Security considerations for any given SMI MIB module will be relevant to any XSD/XML mapping of that MIB module; however, the mapping defined in this document does not itself introduce any new security considerations.

If and when proxies or gateways are developed to convey SNMP management information from SNMP agents to XML-based management applications via XSD/XML mapping of MIB modules based on this specification and its planned siblings, special care will need to be taken to ensure that all applicable SNMP security mechanisms are supported in an appropriate manner yet to be determined.

7. IANA Considerations

In accordance with RFC 3688 [RFC3688], the IANA XML registry has been updated with the following namespace and schema registrations associated with this document:

- o urn:ietf:params:xml:ns:smi:base:1.0
- o urn:ietf:params:xml:schema:base:1.0

7.1. SMI Base Datatypes Namespace Registration

This document registers a URI for the SMI Base Datatypes XML namespace in the IETF XML registry. Following the format in RFC 3688, IANA has made the following registration:

URI: urn:ietf:params:xml:ns:smi:base:1.0

Registration Contact: The IESG.

XML: N/A, the requested URI is an XML namespace.

7.2. SMI Base Datatypes Schema Registration

This document registers a URI for the SMI Base Datatypes XML schema in the IETF XML registry. Following the format in RFC 3688, IANA has made the following registration:

URI: urn:ietf:params:xml:schema:smi:base:1.0

Registration Contact: The IESG.

XML: Section 4 of this document.

8. Acknowledgements

Dave Harrington provided strategic and technical leadership to the team that developed this particular specification. Yan Li did much of the research into existing approaches that was used as a baseline for the recommendations in this particular specification.

This document owes much to "Datatypes for Netconf Data Models" [NETCONF-DATATYPES] and to many other sources (including libsmi and group discussions on the NETCONF mailing lists) developed by those who have researched and published candidate mappings of SMI datatypes to XSD.

Individuals who participated in various discussions of this topic at IETF meetings and on IETF mailing lists include: Ray Atarashi, Yoshifumi Atarashi, Andy Bierman, Sharon Chisholm, Avri Doria, Rob Ennes, Mehmet Ersue, David Harrington, Alfred Hines, Eliot Lear, Chris Lonvick, Faye Ly, Randy Presuhn, Juergen Schoenwaelder, Andrea Westerinen, and Bert Wijnen.

9. References

9.1. Normative References

- [RFC1155] Rose, M. and K. McCloghrie, "Structure and identification of management information for TCP/IP-based internets", STD 16, RFC 1155, May 1990.
- [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.
- [RFC3584] Frye, R., Levi, D., Routhier, S., and B. Wijnen, "Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework", BCP 74, RFC 3584, August 2003.
- [XML] World Wide Web Consortium, "Extensible Markup Language (XML) 1.0", W3C XML, February 1998, <<http://www.w3.org/TR/1998/REC-xml-19980210>>.
- [XMLSchema] World Wide Web Consortium, "XML Schema Part 1: Structures Second Edition", W3C XML Schema, October 2004, <<http://www.w3.org/TR/xmlschema-1/>>.
- [XSDDatatypes] World Wide Web Consortium, "XML Schema Part 2: Datatypes Second Edition", W3C XML Schema, October 2004, <<http://www.w3.org/TR/xmlschema-2/>>.

9.2. Informative References

- [ASN.1] International Organization for Standardization, "Information processing systems - Open Systems Interconnection - Specification of Basic Encoding Rules for Abstract Syntax Notation One (ASN.1)", International Standard 8825, December 1987.
- [NETCONF-DATATYPES] Romascanu, D., "Datatypes for Netconf Data Models", Work in Progress, May 2007.

- [RFC2579] McCloghrie, K., Perkins, D., and J. Schoenwaelder, "Textual Conventions for SMIV2", STD 58, RFC 2579, April 1999.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, January 2004.
- [RFC4001] Daniele, M., Haberman, B., Routhier, S., and J. Schoenwaelder, "Textual Conventions for Internet Network Addresses", RFC 4001, February 2005.

Authors' Addresses

Mark Ellison
Ellison Software Consulting
38 Salem Road
Atkinson, NH 03811
USA

Phone: +1 603-362-9270
EMail: ietf@ellisonsoftware.com

Bob Natale
MITRE
300 Sentinel Drive
6th Floor
Annapolis Junction, MD 20701
USA

Phone: +1 301-617-3008
EMail: rnatale@mitre.org