

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: July 16, 2016

C. Jennings  
Cisco  
Z. Shelby  
ARM  
J. Arkko  
A. Keranen  
Ericsson  
January 13, 2016

**Media Types for Sensor Markup Language (SENML)**  
**draft-jennings-core-senml-04**

**Abstract**

This specification defines media types for representing simple sensor measurements and device parameters in the Sensor Markup Language (SenML). Representations are defined in JavaScript Object Notation (JSON), Concise Binary Object Representation (CBOR), eXtensible Markup Language (XML), and Efficient XML Interchange (EXI), which share the common SenML data model. A simple sensor, such as a temperature sensor, could use this media type in protocols such as HTTP or CoAP to transport the measurements of the sensor or to be configured.

**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 July 16, 2016.

**Copyright Notice**

Copyright (c) 2016 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

<a href="#">1.</a>	Overview . . . . .	<a href="#">3</a>
<a href="#">2.</a>	Requirements and Design Goals . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Terminology . . . . .	<a href="#">5</a>
<a href="#">4.</a>	Semantics . . . . .	<a href="#">5</a>
<a href="#">5.</a>	Associating Meta-data . . . . .	<a href="#">7</a>
<a href="#">6.</a>	JSON Representation (application/senml+json) . . . . .	<a href="#">8</a>
<a href="#">6.1.</a>	Examples . . . . .	<a href="#">9</a>
<a href="#">6.1.1.</a>	Single Datapoint . . . . .	<a href="#">9</a>
<a href="#">6.1.2.</a>	Multiple Datapoints . . . . .	<a href="#">9</a>
<a href="#">6.1.3.</a>	Multiple Measurements . . . . .	<a href="#">10</a>
<a href="#">6.1.4.</a>	Collection of Resources . . . . .	<a href="#">11</a>
<a href="#">7.</a>	CBOR Representation (application/senml+cbor) . . . . .	<a href="#">11</a>
<a href="#">8.</a>	XML Representation (application/senml+xml) . . . . .	<a href="#">12</a>
<a href="#">9.</a>	EXI Representation (application/senml-exi) . . . . .	<a href="#">14</a>
<a href="#">10.</a>	Usage Considerations . . . . .	<a href="#">16</a>
<a href="#">11.</a>	IANA Considerations . . . . .	<a href="#">17</a>
<a href="#">11.1.</a>	Units Registry . . . . .	<a href="#">18</a>
<a href="#">11.2.</a>	Media Type Registration . . . . .	<a href="#">20</a>
<a href="#">11.2.1.</a>	senml+json Media Type Registration . . . . .	<a href="#">20</a>
<a href="#">11.2.2.</a>	senml+cbor Media Type Registration . . . . .	<a href="#">22</a>
<a href="#">11.2.3.</a>	senml+xml Media Type Registration . . . . .	<a href="#">22</a>
<a href="#">11.2.4.</a>	senml-exi Media Type Registration . . . . .	<a href="#">23</a>
<a href="#">11.3.</a>	XML Namespace Registration . . . . .	<a href="#">24</a>
<a href="#">11.4.</a>	CoAP Content-Format Registration . . . . .	<a href="#">24</a>
<a href="#">12.</a>	Security Considerations . . . . .	<a href="#">25</a>
<a href="#">13.</a>	Privacy Considerations . . . . .	<a href="#">25</a>
<a href="#">14.</a>	Acknowledgement . . . . .	<a href="#">25</a>
<a href="#">15.</a>	References . . . . .	<a href="#">25</a>
<a href="#">15.1.</a>	Normative References . . . . .	<a href="#">25</a>
<a href="#">15.2.</a>	Informative References . . . . .	<a href="#">27</a>
	Authors' Addresses . . . . .	<a href="#">27</a>



## 1. Overview

Connecting sensors to the internet is not new, and there have been many protocols designed to facilitate it. This specification defines new media types for carrying simple sensor information in a protocol such as HTTP or CoAP called the Sensor Markup Language (SenML). This format was designed so that processors with very limited capabilities could easily encode a sensor measurement into the media type, while at the same time a server parsing the data could relatively efficiently collect a large number of sensor measurements. The markup language can be used for a variety of data flow models, most notably data feeds pushed from a sensor to a collector, and the web resource model where the sensor is requested as a resource representation (e.g., "GET /sensor/temperature").

There are many types of more complex measurements and measurements that this media type would not be suitable for. SenML strikes a balance between having some information about the sensor carried with the sensor data so that the data is self describing but it also tries to make that a fairly minimal set of auxiliary information for efficiency reason. Other information about the sensor can be discovered by other methods such as using the CoRE Link Format [[RFC6690](#)].

SenML is defined by a data model for measurements and simple meta-data about measurements and devices. The data is structured as a single array that contains a series of SenML Records which can each contain attributes such as an unique identifier for the sensor, the time the measurement was made, the unit the measurement is in, and the current value of the sensor. Serializations for this data model are defined for JSON [[RFC7159](#)], CBOR [[RFC7049](#)], XML, and Efficient XML Interchange (EXI) [[W3C.REC-exi-20110310](#)].

For example, the following shows a measurement from a temperature gauge encoded in the JSON syntax.

```
[{ "n": "urn:dev:ow:10e2073a01080063", "v":23.1, "u":"Cel" }]
```

In the example above, the array has a single SenML record with a measurement for a sensor named "urn:dev:ow:10e2073a01080063" with a current value of 23.5 degrees Celsius.

## 2. Requirements and Design Goals

The design goal is to be able to send simple sensor measurements in small packets on mesh networks from large numbers of constrained devices. Keeping the total size of payload under 80 bytes makes this easy to use on a wireless mesh network. It is always difficult to



define what small code is, but there is a desire to be able to implement this in roughly 1 KB of flash on a 8 bit microprocessor. Experience with Google power meter and large scale deployments has indicated that the solution needs to support allowing multiple measurements to be batched into a single HTTP or CoAP request. This "batch" upload capability allows the server side to efficiently support a large number of devices. It also conveniently supports batch transfers from proxies and storage devices, even in situations where the sensor itself sends just a single data item at a time. The multiple measurements could be from multiple related sensors or from the same sensor but at different times.

The basic design is an array with a series of measurements. The following example shows two measurements made at different times. The value of the measurements are in the "v" tag, the time of the measurement is in the "t" while the "n" has the unique sensor name and unit is carried in the "u".

```
[
  { "n": "urn:dev:ow:10e2073a01080063",
    "t": 1276020076, "v":23.5, "u":"Cel" },
  { "n": "urn:dev:ow:10e2073a01080063",
    "t": 1276020091, "v":23.6, "u":"Cel" }
]
```

To keep the messages small, it does not make sense to repeat the n in each SenML Record so there is a concept of a Base Name which is simply a string that is prepended to the Name field of any elements that record or any records that follow it and don't contain a Base Name. So a more compact form of the example above is the following.

```
[
  { "bn": "urn:dev:ow:10e2073a01080063",
    "t": 1276020076, "v":23.5, "u":"Cel" },
  { "t": 1276020091, "v":23.6, "u":"Cel" }
]
```

In the above example the Base Name is in the "bn" tag and the "n" tags in each Record are the empty string so they are omitted. The Base Name also could be put in a separate Record such as the following example.

```
[
  { "bn": "urn:dev:ow:10e2073a01080063" },
  { "t": 1276020076, "v":23.5, "u":"Cel" },
  { "t": 1276020091, "v":23.6, "u":"Cel" }
]
```



Some devices have accurate time while others do not so SenML supports absolute and relative times. Time is represented in floating point as second and values greater than zero represent an absolute time relative to the unix epoch while values of 0 or less represent a relative time in the past from the current time. A simple sensor with no absolute wall clock time might take a measurement every second and batch up 60 of them then send it to a server. It would include the relative time the measurement was made to the time the batch was send in the SenML. The server might have accurate NTP time and use the time it received the data, and the relative offset, to replace the times in the SenML with absolute times before saving the SenML in a document database.

### **3. Terminology**

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

### **4. Semantics**

Each SenML representation carries a single array that represents a set of measurements and/or parameters. This array contains a series of object with several optional attributes described below:

**Base Name:** This is a string that is prepended to the names found in the entries. This attribute is optional. This applies to the entries in the Record containing the Base Name and all entries that come in Records after it until there is a Record with a new Base Name that replaces this one.

**Base Time:** A base time that is added to the time found in an entry. This attribute is optional. This applies to the entries in the Record containing the Base Time and all entries that come in Records after it until there is a Record with a new Base Time that replaces this one.

**Base Unit:** A base unit that is assumed for all entries, unless otherwise indicated. This attribute is optional. If a record does not a unit value, then the base unit is used otherwise the value of found in the Unit is used. Note the Base Unit is not prepended to the Unit. This applies to the entries in the Record containing the Base Unit and all entries that come in Records after it until there is a Record with a new Base Unit that replaces this one.





**Version:** Version number of media type format. This attribute is optional positive integer and defaults to 3 if not present. If this value is present at all, it SHOULD only be used in the first Record in the SenML Stream array.

**Name:** Name of the sensor or parameter. When appended to the Base Name attribute, this must result in a globally unique identifier for the resource. The name is optional, if the Base Name is present. If the name is missing, Base Name must uniquely identify the resource. This can be used to represent a large array of measurements from the same sensor without having to repeat its identifier on every measurement.

**Unit:** Units for a measurement value. Optional. If the Record has not Unit, the Base Unit is used as the Unit.

**Value** Value of the entry. Optional if a Sum value is present, otherwise required. Values are represented using three basic data types, Floating point numbers ("v" field for "Value"), Booleans ("vb" for "Boolean Value") and Strings ("vs" for "String Value"). Exactly one of these three fields MUST appear.

**Sum:** Integrated sum of the values over time. Optional. This attribute is in the units specified in the Unit value multiplied by seconds.

**Time:** Time when value was recorded. Optional.

**Update Time:** A time in seconds that represents the maximum time before this sensor will provide an updated reading for a measurement. This can be used to detect the failure of sensors or communications path from the sensor. Optional.

The SenML format can be extended with further custom attributes.  
TODO - describe what extensions are possible and how to do them.

Systems reading one of the objects MUST check for the Version attribute. If this value is a version number larger than the version which the system understands, the system SHOULD NOT use this object. This allows the version number to indicate that the object contains mandatory to understand attributes. New version numbers can only be defined in an RFC that updates this specification or its successors.

The Name value is concatenated to the Base Name value to get the name of the sensor. The resulting name needs to uniquely identify and differentiate the sensor from all others. If the object is a representation resulting from the request of a URI [[RFC3986](#)], then in the absence of the Base Name attribute, this URI is used as the



default value of Base Name. Thus in this case the Name field needs to be unique for that URI, for example an index or subresource name of sensors handled by the URI.

Alternatively, for objects not related to a URI, a unique name is required. In any case, it is RECOMMENDED that the full names are represented as URIs or URNs [[RFC2141](#)]. One way to create a unique name is to include a EUI-48 or EUI-64 identifier (a MAC address) or some other bit string that has guaranteed uniqueness (such as a 1-wire address) that is assigned to the device. Some of the examples in this draft use the device URN type as specified in [[I-D.arkko-core-dev-urn](#)]. UUIDs [[RFC4122](#)] are another way to generate a unique name.

The resulting concatenated name MUST consist only of characters out of the set "A" to "Z", "a" to "z", "0" to "9", "-", ":", ".", or "\_" and it MUST start with a character out of the set "A" to "Z", "a" to "z", or "0" to "9". This restricted character set was chosen so that these names can be directly used as in other types of URI including segments of an HTTP path with no special encoding and can be directly used in many databases and analytic systems.. [[RFC5952](#)] contains advice on encoding an IPv6 address in a name.

If either the Base Time or Time value is missing, the missing attribute is considered to have a value of zero. The Base Time and Time values are added together to get the time of measurement. A time of zero indicates that the sensor does not know the absolute time and the measurement was made roughly "now". A negative value is used to indicate seconds in the past from roughly "now". A positive value is used to indicate the number of seconds, excluding leap seconds, since the start of the year 1970 in UTC.

Representing the statistical characteristics of measurements, such as accuracy, can be very complex. Future specification may add new attributes to provide better information about the statistical properties of the measurement.

## **5. Associating Meta-data**

SenML is designed to carry the minimum dynamic information about measurements, and for efficiency reasons does not carry significant static meta-data about the device, object or sensors. Instead, it is assumed that this meta-data is carried out of band. For web resources using SenML representations, this meta-data can be made available using the CoRE Link Format [[RFC6690](#)]. The most obvious use of this link format is to describe that a resource is available in a SenML format in the first place. The relevant media type indicator is included in the Content-Type (ct=) attribute.



## 6. JSON Representation (application/senml+json)

Record attributes:

SenML	JSON	Type
Base Name	bn	String
Base Time	bt	Number
Base Unit	bu	Number
Version	ver	Number
Name	n	String
Unit	u	String
Value	v	Floating point
String Value	vs	String
Boolean Value	vb	Boolean
Value Sum	s	Floating point
Time	t	Number
Update Time	ut	Number

All of the data is UTF-8, but since this is for machine to machine communications on constrained systems, only characters with code points between U+0001 and U+007F are allowed which corresponds to the ASCII [\[RFC0020\]](#) subset of UTF-8 with the exception of characters found in the String Value.

Characters in the String Value are encoded TODO. Open Issue How.

The root content consists of an array with and JSON object for each SenML Record.

The objects MAY contain a "bn" attribute with a value of type string. The object MAY contain a "bt" attribute with a value of type number. The object MAY contain a "bu" attribute with a value of type string. The object MAY contain a "ver" attribute with a value of type number. The object MAY contain other attribute value pairs.

The objects MAY include the "n", "u", and "vs" attributes are of type string, the "t" and "ut" attributes are of type number, the "vb" attribute is of type boolean, and the "v" and "s" attributes are of type floating point for the SenML attributes defined in the table above. All the attributes are optional, but as specified in [Section 4](#), one of the "v", "vs", or "vb" attributes MUST appear unless the "s" attribute is also present in Records that represent a measurement. The "v", and "vs", and "vb" attributes MUST NOT appear together in the same object.



Systems receiving measurements MUST be able to process the range of floating point numbers that are representable as an IEEE double-precision floating-point numbers [[IEEE.754.1985](#)]. The number of significant digits in any measurement is not relevant, so a reading of 1.1 has exactly the same semantic meaning as 1.10. If the value has an exponent, the "e" MUST be in lower case. The mantissa SHOULD be less than 19 characters long and the exponent SHOULD be less than 5 characters long. This allows time values to have better than micro second precision over the next 100 years.

## **[6.1.](#) Examples**

### **[6.1.1.](#) Single Datapoint**

The following shows a temperature reading taken approximately "now" by a 1-wire sensor device that was assigned the unique 1-wire address of 10e2073a01080063:

```
[{ "n": "urn:dev:ow:10e2073a01080063", "v":23.1, "u":"Cel" }]
```

### **[6.1.2.](#) Multiple Datapoints**

The following example shows voltage and current now, i.e., at an unspecified time. The device has an EUI-64 MAC address of 0024beffffe804ff1.

```
[{"bn": "urn:dev:mac:0024beffffe804ff1/"},
 { "n": "voltage", "t": 0, "u": "V", "v": 120.1 },
 { "n": "current", "t": 0, "u": "A", "v": 1.2 }
]
```

The next example is similar to the above one, but shows current at Tue Jun 8 18:01:16 UTC 2010 and at each second for the previous 5 seconds.

```
[{"bn": "urn:dev:mac:0024beffffe804ff1/",
 "bt": 1276020076,
 "bu": "A",
 "ver": 3},
 { "n": "voltage", "u": "V", "v": 120.1 },
 { "n": "current", "t": -5, "v": 1.2 },
 { "n": "current", "t": -4, "v": 1.30 },
 { "n": "current", "t": -3, "v": 0.14e1 },
 { "n": "current", "t": -2, "v": 1.5 },
 { "n": "current", "t": -1, "v": 1.6 },
 { "n": "current", "t": 0, "v": 1.7 }
]
```





Note that in some usage scenarios of SenML the implementations MAY store or transmit SenML in a stream-like fashion, where data is collected over time and continuously added to the object. This mode of operation is optional, but systems or protocols using SenML in this fashion MUST specify that they are doing this. SenML defines a separate mime type (TODO) to indicate Sensor Streaming Markup Language (SensML) for this usage. In this situation the SensML stream can be sent and received in a partial fashion, i.e., a measurement entry can be read as soon as the SenML Record is received and not have to wait for the full SenML Stream to be complete.

For instance, the following stream of measurements may be sent via a long lived HTTP POST from the producer of a SensML to the consumer of that, and each measurement object may be reported at the time it measured:

```
[ {"bn": "http://[2001:db8::1]",
  "bt": 1320067464,
  "bu": "%RH"},
  { "v": 21.2, "t": 0 },
  { "v": 21.3, "t": 10 },
  { "v": 21.4, "t": 20 },
  { "v": 21.4, "t": 30 },
  { "v": 21.5, "t": 40 },
  { "v": 21.5, "t": 50 },
  { "v": 21.5, "t": 60 },
  { "v": 21.6, "t": 70 },
  { "v": 21.7, "t": 80 },
  { "v": 21.5, "t": 90 },
  ...
```

### **6.1.3. Multiple Measurements**

The following example shows humidity measurements from a mobile device with an IPv6 address 2001:db8::1, starting at Mon Oct 31 13:24:24 UTC 2011. The device also provides position data, which is provided in the same measurement or parameter array as separate entries. Note time is used to for correlating data that belongs together, e.g., a measurement and a parameter associated with it. Finally, the device also reports extra data about its battery status at a separate time.



```
[{"bn": "http://[2001:db8::1]",
  "bt": 1320067464,
  "bu": "%RH"},
  { "v": 20.0, "t": 0 },
  { "v": 24.30621, "u": "lon", "t": 0 },
  { "v": 60.07965, "u": "lat", "t": 0 },
  { "v": 20.3, "t": 60 },
  { "v": 24.30622, "u": "lon", "t": 60 },
  { "v": 60.07965, "u": "lat", "t": 60 },
  { "v": 20.7, "t": 120 },
  { "v": 24.30623, "u": "lon", "t": 120 },
  { "v": 60.07966, "u": "lat", "t": 120 },
  { "v": 98.0, "u": "%EL", "t": 150 },
  { "v": 21.2, "t": 180 },
  { "v": 24.30628, "u": "lon", "t": 180 },
  { "v": 60.07967, "u": "lat", "t": 180 }
]
```

#### [6.1.4.](#) Collection of Resources

The following example shows how to query one device that can provide multiple measurements. The example assumes that a client has fetched information from a device at 2001:db8::2 by performing a GET operation on `http://[2001:db8::2]` at Mon Oct 31 16:27:09 UTC 2011, and has gotten two separate values as a result, a temperature and humidity measurement.

```
[{"bn": "http://[2001:db8::2]/",
  "bt": 1320078429,
  "ver": 3},
  { "n": "temperature", "v": 27.2, "u": "Cel" },
  { "n": "humidity", "v": 80, "u": "%RH" }
]
```

## [7.](#) CBOR Representation (`application/senml+cbor`)

The CBOR [\[RFC7049\]](#) representation is equivalent to the JSON representation, with the following changes:

- o For compactness, the CBOR representation uses integers for the map keys defined in Table 1. This table is conclusive, i.e., there is no intention to define any additional integer map keys; any extensions will use string map keys.
- o For JSON Numbers, the CBOR representation can use integers, floating point numbers, or decimal fractions (CBOR Tag 4); the common limitations of JSON implementations are not relevant for



these. For the version number, however, only an unsigned integer is allowed.

Name	JSON label	CBOR label
Version	ver	-1
Base Name	bn	-2
Base Time	bt	-3
Base Units	bu	-4
Name	n	0
Units	u	1
Value	v	2
String Value	vs	3
Boolean Value	vb	4
Value Sum	s	5
Time	t	6
Update Time	ut	7

Table 1: CBOR representation: integers for map keys

## 8. XML Representation (application/senml+xml)

A SenML Stream can also be represented in XML format as defined in this section. The following example shows an XML example for the same sensor measurement as in [Section 6.1.2](#).

```
<senmls xmlns="urn:ietf:params:xml:ns:senml">
  <senml bn="urn:dev:mac:0024beffffe804ff1/" bt="1276020076" bu="A"
    ver="3"></senml>
  <senml n="voltage" u="V" v="120.1"></senml>
  <senml n="current" t="-5" v="1.2"></senml>
  <senml n="current" t="-4" v="1.3"></senml>
  <senml n="current" t="-3" v="1.4"></senml>
  <senml n="current" t="-2" v="1.5"></senml>
  <senml n="current" t="-1" v="1.6"></senml>
  <senml n="current" v="1.7"></senml>
</senmls>
```

TODO - tag names in examples are wrong

The SenML Stream is represented as a `sensml` tag that contains a series of `senml` tags for each SenML Record. The SenML Fields are represents as XML attributes. The following table shows the mapping the SenML Field names to the attribute used on the XML `senml` tag.



SenML Field	XML	Type
Base Name	bn	string
Base Time	bt	int
Base Unit	bu	int
Version	ver	int
Name	n	string
Unit	u	string
Value	v	float
String Value	vs	string
Boolean Value	vb	boolean
Value Sum	s	float
Time	t	int
Update Time	ut	int

The RelaxNG schema for the XML is:

```
default namespace = "urn:ietf:params:xml:ns:senml"
namespace rng = "http://relaxng.org/ns/structure/1.0"
```

```
senml = element senml {
  attribute bn { xsd:string }?,
  attribute bt { xsd:int }?,
  attribute bu { xsd:string }?,
  attribute n { xsd:string }?,
  attribute s { xsd:float }?,
  attribute t { xsd:int }?,
  attribute u { xsd:string }?,
  attribute ut { xsd:int }?,
  attribute v { xsd:float }?,
  attribute vb { xsd:boolean }?,
  attribute ver { xsd:int }?,
  attribute vs { xsd:string }?
}
```

```
senmls =
  element senmls {
    senml+
  }
```

```
start = senmls
```





## **9. EXI Representation (application/senml-exi)**

For efficient transmission of SenML over e.g. a constrained network, Efficient XML Interchange (EXI) can be used. This encodes the XML Schema structure of SenML into binary tags and values rather than ASCII text. An EXI representation of SenML SHOULD be made using the strict schema-mode of EXI. This mode however does not allow tag extensions to the schema, and therefore any extensions will be lost in the encoding. For uses where extensions need to be preserved in EXI, the non-strict schema mode of EXI MAY be used.

The EXI header option MUST be included. An EXI schemaID options MUST be set to the value of "a" indicating the scheme provided in this specification. Future revisions to the schema can change this schemaID to allow for backwards compatibility. When the data will be transported over CoAP or HTTP, an EXI Cookie SHOULD NOT be used as it simply makes things larger and is redundant to information provided in the Content-Type header.

TODO - examples are probably have the wrong setting the schemaID

The following is the XSD Schema to be used for strict schema guided EXI processing. It is generated from the RelaxNG.



```

<?xml version="1.0" encoding="utf-8"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
  elementFormDefault="qualified"
  targetNamespace="urn:ietf:params:xml:ns:senml"
  xmlns:ns1="urn:ietf:params:xml:ns:senml">
  <xs:element name="senml">
    <xs:complexType>
      <xs:attribute name="bn" type="xs:string" />
      <xs:attribute name="bt" type="xs:int" />
      <xs:attribute name="bu" type="xs:string" />
      <xs:attribute name="n" type="xs:string" />
      <xs:attribute name="s" type="xs:float" />
      <xs:attribute name="t" type="xs:int" />
      <xs:attribute name="u" type="xs:string" />
      <xs:attribute name="ut" type="xs:int" />
      <xs:attribute name="v" type="xs:float" />
      <xs:attribute name="vb" type="xs:boolean" />
      <xs:attribute name="ver" type="xs:int" />
      <xs:attribute name="vs" type="xs:string" />
    </xs:complexType>
  </xs:element>
  <xs:element name="senmls">
    <xs:complexType>
      <xs:sequence>
        <xs:element maxOccurs="unbounded" ref="ns1:senml" />
      </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:schema>

```

The following shows a hexdump of the EXI produced from encoding the following XML example. Note this example is the same information as the first example in [Section 6.1.2](#) in JSON format.

```

<senmls xmlns="urn:ietf:params:xml:ns:senml">
  <senml bn="urn:dev:mac:0024beffffe804ff1/"></senml>
  <senml n="voltage" u="V" v="120.1"></senml>
  <senml n="current" u="A" v="1.2"></senml>
</senmls>

```

TODO - replace ex8 with ex2 as they are the same. Replace ex9 with ex1. Simplify rest of examples.

Which compresses to the following displayed in hexdump:



```

00000000 a0 30 41 cd 95 b9 b5 b0 cc b9 9d 95 b8 b9 e1 cd
00000010 90 80 fb ab 93 71 d3 23 2b b1 d3 6b 0b 19 d1 81
00000020 81 91 a3 13 2b 33 33 33 29 c1 81 a3 33 31 89 7d
00000030 8c 25 d9 bd b1 d1 85 9d 94 80 d5 8a c4 26 01 8c
00000040 25 8d d5 c9 c9 95 b9 d0 80 d0 48 32 01 c0
0000004e

```

The above example used the bit packed form of EXI but it is also possible to use a byte packed form of EXI which can makes it easier for a simple sensor to produce valid EXI without really implementing EXI. Consider the example of a temperature sensor that produces a value in tenths of degrees Celsius over a range of 0.0 to 55.0. It would produce an XML SenML file such as:

```

<senmls xmlns="urn:ietf:params:xml:ns:senml">
  <senml n="urn:dev:ow:10e2073a01080063" u="Cel" v="23.1"></senml>
</senmls>

```

The compressed form, using the byte alignment option of EXI, for the above XML is the following:

```

00 a0 00 48 82 0e 6c ad cd ad 86 65 cc ec ad c5 cf |...H..l....e....|
10 0e 6c 80 01 03 1d 75 72 6e 3a 64 65 76 3a 6f 77 |.l....urn:dev:ow|
20 3a 31 30 65 32 30 37 33 61 30 31 30 38 30 30 36 |:10e2073a0108006|
30 33 02 05 43 65 6c 01 00 e7 01 01 00 03 01      |3..Cel.....|
3e

```

A small temperature sensor devices that only generates this one EXI file does not really need an full EXI implementation. It can simple hard code the output replacing the one wire device ID starting at byte 0x16 and going to byte 0x31 with it's device ID, and replacing the value "0xe7 0x01" at location 0x38 to 0x39 with the current temperature. The EXI Specification [[W3C.REC-exi-20110310](#)] contains the full information 'on how floating point numbers are represented, but for the purpose of this sensor, the temperature can be converted to an integer in tenths of degrees (231 in this example). EXI stores 7 bits of the integer in each byte with the top bit set to one if there are further bytes. So the first bytes at is set to low 7 bits of the integer temperature in tenths of degrees plus 0x80. In this example  $231 \& 0x7F + 0x80 = 0xE7$ . The second byte is set to the integer temperature in tenths of degrees right shifted 7 bits. In this example  $231 \gg 7 = 0x01$ .

## 10. Usage Considerations

The measurements support sending both the current value of a sensor as well as the an integrated sum. For many types of measurements, the sum is more useful than the current value. For example, an



electrical meter that measures the energy a given computer uses will typically want to measure the cumulative amount of energy used. This is less prone to error than reporting the power each second and trying to have something on the server side sum together all the power measurements. If the network between the sensor and the meter goes down over some period of time, when it comes back up, the cumulative sum helps reflect what happened while the network was down. A meter like this would typically report a measurement with the units set to watts, but it would put the sum of energy used in the "s" attribute of the measurement. It might optionally include the current power in the "v" attribute.

While the benefit of using the integrated sum is fairly clear for measurements like power and energy, it is less obvious for something like temperature. Reporting the sum of the temperature makes it easy to compute averages even when the individual temperature values are not reported frequently enough to compute accurate averages. Implementors are encouraged to report the cumulative sum as well as the raw value of a given sensor.

Applications that use the cumulative sum values need to understand they are very loosely defined by this specification, and depending on the particular sensor implementation may behave in unexpected ways. Applications should be able to deal with the following issues:

1. Many sensors will allow the cumulative sums to "wrap" back to zero after the value gets sufficiently large.
2. Some sensors will reset the cumulative sum back to zero when the device is reset, loses power, or is replaced with a different sensor.
3. Applications cannot make assumptions about when the device started accumulating values into the sum.

Typically applications can make some assumptions about specific sensors that will allow them to deal with these problems. A common assumption is that for sensors whose measurement values are always positive, the sum should never get smaller; so if the sum does get smaller, the application will know that one of the situations listed above has happened.

## **11. IANA Considerations**

Note to RFC Editor: Please replace all occurrences of "RFC-AAAA" with the RFC number of this specification.





### 11.1. Units Registry

IANA will create a registry of unit symbols. The primary purpose of this registry is to make sure that symbols uniquely map to give type of measurement. Definitions for many of these units can be found in location such as [\[NIST811\]](#) and [\[BIPM\]](#).

Symbol	Description	Type	Reference
m	meter	float	RFC-AAAA
kg	kilogram	float	RFC-AAAA
s	second	float	RFC-AAAA
A	ampere	float	RFC-AAAA
K	kelvin	float	RFC-AAAA
cd	candela	float	RFC-AAAA
mol	mole	float	RFC-AAAA
Hz	hertz	float	RFC-AAAA
rad	radian	float	RFC-AAAA
sr	steradian	float	RFC-AAAA
N	newton	float	RFC-AAAA
Pa	pascal	float	RFC-AAAA
J	joule	float	RFC-AAAA
W	watt	float	RFC-AAAA
C	coulomb	float	RFC-AAAA
V	volt	float	RFC-AAAA
F	farad	float	RFC-AAAA
Ohm	ohm	float	RFC-AAAA
S	siemens	float	RFC-AAAA
Wb	weber	float	RFC-AAAA
T	tesla	float	RFC-AAAA
H	henry	float	RFC-AAAA
Cel	degrees Celsius	float	RFC-AAAA
lm	lumen	float	RFC-AAAA
lx	lux	float	RFC-AAAA
Bq	becquerel	float	RFC-AAAA
Gy	gray	float	RFC-AAAA
Sv	sievert	float	RFC-AAAA
kat	katal	float	RFC-AAAA
pH	pH acidity	float	RFC-AAAA
%	Value of a switch (note 1)	float	RFC-AAAA
count	counter value	float	RFC-AAAA
%RH	Relative Humidity	float	RFC-AAAA
m2	area	float	RFC-AAAA
l	volume in liters	float	RFC-AAAA
m/s	velocity	float	RFC-AAAA
m/s2	acceleration	float	RFC-AAAA
l/s	flow rate in liters per second	float	RFC-AAAA



	W/m2	irradiance		float	RFC-AAAA	
	cd/m2	luminance		float	RFC-AAAA	
	Bspl	bel sound pressure level		float	RFC-AAAA	
	bit/s	bits per second		float	RFC-AAAA	
	lat	degrees latitude (note 2)		float	RFC-AAAA	
	lon	degrees longitude (note 2)		float	RFC-AAAA	
	%EL	remaining battery energy level in		float	RFC-AAAA	
		percents				
	EL	remaining battery energy level in		float	RFC-AAAA	
		seconds				
	beat/m	Heart rate in beats per minute		float	RFC-AAAA	
	beats	Cumulative number of heart beats		float	RFC-AAAA	
+-----+-----+-----+-----+-----+-----+						

Table 2

- o Note 1: A value of 0.0 indicates the switch is off while 1.0 indicates on and 0.5 would be half on.
- o Note 2: Assumed to be in WGS84 unless another reference frame is known for the sensor.

New entries can be added to the registration by either Expert Review or IESG Approval as defined in [[RFC5226](#)]. Experts should exercise their own good judgment but need to consider the following guidelines:

1. There needs to be a real and compelling use for any new unit to be added.
2. Units should define the semantic information and be chosen carefully. Implementors need to remember that the same word may be used in different real-life contexts. For example, degrees when measuring latitude have no semantic relation to degrees when measuring temperature; thus two different units are needed.
3. These measurements are produced by computers for consumption by computers. The principle is that conversion has to be easily be done when both reading and writing the media type. The value of a single canonical representation outweighs the convenience of easy human representations or loss of precision in a conversion.
4. Use of SI prefixes such as "k" before the unit is not allowed. Instead one can represent the value using scientific notation such a 1.2e3.
5. For a given type of measurement, there will only be one unit type defined. So for length, meters are defined and other



lengths such as mile, foot, light year are not allowed. For most cases, the SI unit is preferred.

6. Symbol names that could be easily confused with existing common units or units combined with prefixes should be avoided. For example, selecting a unit name of "mph" to indicate something that had nothing to do with velocity would be a bad choice, as "mph" is commonly used to mean miles per hour.
7. The following should not be used because they are common SI prefixes: Y, Z, E, P, T, G, M, k, h, da, d, c, n, u, p, f, a, z, y, Ki, Mi, Gi, Ti, Pi, Ei, Zi, Yi.
8. The following units should not be used as they are commonly used to represent other measurements: Ky, Gal, dyn, etg, P, St, Mx, G, Oe, Gb, sb, Lmb, ph, Ci, R, RAD, REM, gal, bbl, qt, degF, Cal, BTU, HP, pH, B/s, psi, Torr, atm, at, bar, kWh.
9. The unit names are case sensitive and the correct case needs to be used, but symbols that differ only in case should not be allocated.
10. A number after a unit typically indicates the previous unit raised to that power, and the / indicates that the units that follow are the reciprocal. A unit should have only one / in the name.
11. A good list of common units can be found in the Unified Code for Units of Measure [[UCUM](#)].

## **[11.2.](#) Media Type Registration**

The following registrations are done following the procedure specified in [[RFC6838](#)] and [[RFC7303](#)].

### **[11.2.1.](#) senml+json Media Type Registration**

Type name: application

Subtype name: senml+json and sensml+json

Required parameters: none

Optional parameters: none

Encoding considerations: Must be encoded as using a subset of the encoding allowed in [[RFC7159](#)]. See RFC-AAAA for details. This simplifies implementation of very simple system and does not impose



any significant limitations as all this data is meant for machine to machine communications and is not meant to be human readable.

Security considerations: Sensor data can contain a wide range of information ranging from information that is very public, such the outside temperature in a given city, to very private information that requires integrity and confidentiality protection, such as patient health information. This format does not provide any security and instead relies on the transport protocol that carries it to provide security. Given applications need to look at the overall context of how this media type will be used to decide if the security is adequate.

Interoperability considerations: Applications should ignore any JSON key value pairs that they do not understand. This allows backwards compatibility extensions to this specification. The "ver" field can be used to ensure the receiver supports a minimal level of functionality needed by the creator of the JSON object.

Published specification: RFC-AAAA

Applications that use this media type: The type is used by systems that report electrical power usage and environmental information such as temperature and humidity. It can be used for a wide range of sensor reporting systems.

Additional information:

Magic number(s): none

File extension(s): senml

Macintosh file type code(s): none

Person & email address to contact for further information: Cullen Jennings <fluffy@iii.ca>

Intended usage: COMMON

Restrictions on usage: None

Author: Cullen Jennings <fluffy@iii.ca>

Change controller: IESG





### **11.2.2. senml+cbor Media Type Registration**

Type name: application

Subtype name: senml+cbor

Required parameters: none

Optional parameters: none

Encoding considerations: TBD

Security considerations: TBD

Interoperability considerations: TBD

Published specification: RFC-AAAA

Applications that use this media type: The type is used by systems that report electrical power usage and environmental information such as temperature and humidity. It can be used for a wide range of sensor reporting systems.

Additional information:

Magic number(s): none

File extension(s): senml

Macintosh file type code(s): none

Person & email address to contact for further information: Cullen Jennings <fluffy@iii.ca>

Intended usage: COMMON

Restrictions on usage: None

Author: Cullen Jennings <fluffy@iii.ca>

Change controller: IESG

### **11.2.3. senml+xml Media Type Registration**

Type name: application

Subtype name: senml+xml and sensml+xml



Required parameters: none

Optional parameters: none

Encoding considerations: TBD

Security considerations: TBD

Interoperability considerations: TBD

Published specification: RFC-AAAA

Applications that use this media type: TBD

Additional information:

Magic number(s): none

File extension(s): senml

Macintosh file type code(s): none

Person & email address to contact for further information: Cullen Jennings <fluffy@iii.ca>

Intended usage: COMMON

Restrictions on usage: None

Author: Cullen Jennings <fluffy@iii.ca>

Change controller: IESG

#### **11.2.4. senml-exi Media Type Registration**

Type name: application

Subtype name: senml-exi

Required parameters: none

Optional parameters: none

Encoding considerations: TBD

Security considerations: TBD

Interoperability considerations: TBD



Published specification: RFC-AAAA

Applications that use this media type: TBD

Additional information:

Magic number(s): none

File extension(s): senml

Macintosh file type code(s): none

Person & email address to contact for further information: Cullen Jennings <fluffy@iii.ca>

Intended usage: COMMON

Restrictions on usage: None

Author: Cullen Jennings <fluffy@iii.ca>

Change controller: IESG

### **11.3. XML Namespace Registration**

This document registers the following XML namespaces in the IETF XML registry defined in [[RFC3688](#)].

URI: urn:ietf:params:xml:ns:senml

Registrant Contact: The IESG.

XML: N/A, the requested URIs are XML namespaces

### **11.4. CoAP Content-Format Registration**

IANA is requested to assign CoAP Content-Format IDs for the SenML media types in the "CoAP Content-Formats" sub-registry, within the "CoRE Parameters" registry [[RFC7252](#)]. All IDs are assigned from the "Expert Review" (0-255) range. The assigned IDs are show in Table 3.



Media type	ID
application/senml+json	TBD
application/sensml+json	TBD
application/senml+cbor	TBD
application/senml+xml	TBD
application/sensml+xml	TBD
application/senml-exi	TBD

Table 3: CoAP Content-Format IDs

## 12. Security Considerations

See [Section 13](#). Further discussion of security properties can be found in [Section 11.2](#).

## 13. Privacy Considerations

Sensor data can range from information with almost no security considerations, such as the current temperature in a given city, to highly sensitive medical or location data. This specification provides no security protection for the data but is meant to be used inside another container or transport protocol such as S/MIME or HTTP with TLS that can provide integrity, confidentiality, and authentication information about the source of the data.

## 14. Acknowledgement

We would like to thank Lisa Dusseault, Joe Hildebrand, Lyndsay Campbell, Martin Thomson, John Klensin, Bjoern Hoehrmann, Carsten Bormann, and Christian Amsuess for their review comments.

The CBOR Representation text was contributed by Carsten Bormann.

## 15. References

### 15.1. Normative References

- [BIPM] Bureau International des Poids et Mesures, "The International System of Units (SI)", 8th edition, 2006.
- [IEEE.754.1985] Institute of Electrical and Electronics Engineers, "Standard for Binary Floating-Point Arithmetic", IEEE Standard 754, August 1985.





- [NIST811] Thompson, A. and B. Taylor, "Guide for the Use of the International System of Units (SI)", NIST Special Publication 811, 2008.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/[RFC2119](#), March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3688] Mealling, M., "The IETF XML Registry", [BCP 81](#), [RFC 3688](#), DOI 10.17487/RFC3688, January 2004, <<http://www.rfc-editor.org/info/rfc3688>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC6838] Freed, N., Klensin, J., and T. Hansen, "Media Type Specifications and Registration Procedures", [BCP 13](#), [RFC 6838](#), DOI 10.17487/RFC6838, January 2013, <<http://www.rfc-editor.org/info/rfc6838>>.
- [RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", [RFC 7049](#), DOI 10.17487/RFC7049, October 2013, <<http://www.rfc-editor.org/info/rfc7049>>.
- [RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", [RFC 7159](#), DOI 10.17487/RFC7159, March 2014, <<http://www.rfc-editor.org/info/rfc7159>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", [RFC 7252](#), DOI 10.17487/[RFC7252](#), June 2014, <<http://www.rfc-editor.org/info/rfc7252>>.
- [RFC7303] Thompson, H. and C. Lilley, "XML Media Types", [RFC 7303](#), DOI 10.17487/RFC7303, July 2014, <<http://www.rfc-editor.org/info/rfc7303>>.
- [W3C.REC-exi-20110310]  
Schneider, J. and T. Kamiya, "Efficient XML Interchange (EXI) Format 1.0", World Wide Web Consortium Recommendation REC-exi-20110310, March 2011, <<http://www.w3.org/TR/2011/REC-exi-20110310>>.



## 15.2. Informative References

- [I-D.arkko-core-dev-urn]  
Arkko, J., Jennings, C., and Z. Shelby, "Uniform Resource Names for Device Identifiers", [draft-arkko-core-dev-urn-03](#) (work in progress), July 2012.
- [RFC0020] Cerf, V., "ASCII format for network interchange", STD 80, [RFC 20](#), DOI 10.17487/RFC0020, October 1969, <<http://www.rfc-editor.org/info/rfc20>>.
- [RFC2141] Moats, R., "URN Syntax", [RFC 2141](#), DOI 10.17487/RFC2141, May 1997, <<http://www.rfc-editor.org/info/rfc2141>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), DOI 10.17487/RFC3986, January 2005, <<http://www.rfc-editor.org/info/rfc3986>>.
- [RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique IDentifier (UUID) URN Namespace", [RFC 4122](#), DOI 10.17487/RFC4122, July 2005, <<http://www.rfc-editor.org/info/rfc4122>>.
- [RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", [RFC 5952](#), DOI 10.17487/RFC5952, August 2010, <<http://www.rfc-editor.org/info/rfc5952>>.
- [RFC6690] Shelby, Z., "Constrained RESTful Environments (CoRE) Link Format", [RFC 6690](#), DOI 10.17487/RFC6690, August 2012, <<http://www.rfc-editor.org/info/rfc6690>>.
- [UCUM] Schadow, G. and C. McDonald, "The Unified Code for Units of Measure (UCUM)", Regenstrief Institute and Indiana University School of Informatics, 2013, <<http://unitsofmeasure.org/ucum.html>>.

### Authors' Addresses

Cullen Jennings  
Cisco  
400 3rd Avenue SW  
Calgary, AB T2P 4H2  
Canada

Phone: +1 408 421-9990  
Email: fluffy@cisco.com



Zach Shelby  
ARM  
150 Rose Orchard  
San Jose 95134  
USA

Phone: +1-408-203-9434  
Email: zach.shelby@arm.com

Jari Arkko  
Ericsson  
Jorvas 02420  
Finland

Email: jari.arkko@piuha.net

Ari Keranen  
Ericsson  
Jorvas 02420  
Finland

Email: ari.keranen@ericsson.com

