

Workgroup: Delay-Tolerant Networking
Internet-Draft: draft-birrane-dtn-ari-04
Published: 3 January 2024
Intended Status: Standards Track
Expires: 6 July 2024
Authors: E.J. Birrane E.A. Annis B. Sipos
 JHU/APL JHU/APL JHU/APL

DTNMA Application Resource Identifier (ARI)

Abstract

This document defines the structure, format, and features of the naming scheme for the objects defined in the Delay-Tolerant Networking Management Architecture (DTNMA) Application Management Model (AMM), in support of challenged network management solutions described in the DTNMA document.

This document defines the DTNMA Application Resource Identifier (ARI), using a text-form based on the common Uniform Resource Identifier (URI) and a binary-form based on Concise Binary Object Representation (CBOR). These meet the needs for a concise, typed, parameterized, and hierarchically organized set of managed data elements.

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 <https://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 6 July 2024.

Copyright Notice

Copyright (c) 2024 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
 - [1.1. Scope](#)
 - [1.2. Use of ABNF](#)
 - [1.3. Use of CDDL](#)
 - [1.4. Terminology](#)
- [2. ARI Purpose](#)
 - [2.1. Resource Parameterization](#)
 - [2.2. Compressible Structure](#)
 - [2.2.1. Enumerated Path Segments](#)
 - [2.2.2. Relative Paths](#)
 - [2.2.3. Patterning](#)
- [3. ARI Logical Structure](#)
 - [3.1. Names, Enumerations, Comparisons, and Canonicalizations](#)
 - [3.2. Literals](#)
 - [3.3. Object References](#)
 - [3.3.1. Namespace ID](#)
 - [3.3.2. Object Type](#)
 - [3.3.3. Object ID](#)
 - [3.3.4. Parameters](#)
- [4. ARI Text Form](#)
 - [4.1. URIs and Percent Encoding](#)
 - [4.2. Literals](#)
 - [4.2.1. Typed Literal Values](#)
 - [4.2.2. Untyped Literal Values](#)
 - [4.2.3. Preferred Encodings](#)
 - [4.3. Object References](#)
 - [4.4. URI References](#)
- [5. ARI Binary Form](#)
 - [5.1. Intermediate CBOR](#)
 - [5.2. Literals](#)
 - [5.3. Object References](#)
 - [5.4. URI References](#)
- [6. ARI Patterns](#)
 - [6.1. ARI Matching](#)
- [7. Transcoding Considerations](#)
- [8. Interoperability Considerations](#)
- [9. Security Considerations](#)
- [10. IANA Considerations](#)
 - [10.1. URI Schemes Registry](#)
 - [10.2. CBOR Tags Registry](#)

[10.3. DTN Management Architecture Parameters](#)

[11. References](#)

[11.1. Normative References](#)

[11.2. Informative References](#)

[Appendix A. Examples](#)

[A.1. Primitive-Typed Literal](#)

[A.2. Timestamp Literal](#)

[A.3. Semantic-Typed Literal](#)

[A.4. Complex CBOR Literal](#)

[A.5. Non-parameterized Object Reference](#)

[A.6. Parameterized Object Reference](#)

[A.7. Recursive Structure with Percent Encodings](#)

[Authors' Addresses](#)

1. Introduction

The unique limitations of Delay-Tolerant Networking (DTN) transport capabilities [[RFC4838](#)] necessitate increased reliance on individual node behavior. These limitations are considered part of the expected operational environment of the system and, thus, contemporaneous end-to-end data exchange cannot be considered a requirement for successful communication.

The primary DTN transport mechanism, Bundle Protocol version 7 (BPv7) [[RFC9171](#)], standardizes a store-and-forward behavior required to communicate effectively between endpoints that may never co-exist in a single network partition. BPv7 might be deployed in static environments, but the design and operation of BPv7 cannot presume that to be the case.

Similarly, the management of any BPv7 protocol agent (BPA) (or any software reliant upon DTN for its communication) cannot presume to operate in a resourced, connected network. Just as DTN transport must be delay-tolerant, DTN network management must also be delay-tolerant.

The DTN Management Architecture (DTNMA) [[I-D.ietf-dtn-dtnma](#)] outlines an architecture that achieves this result through the self-management of a DTN node as configured by one or more remote managers in an asynchronous and open-loop system. An important part of this architecture is the definition of a conceptual data schema for defining resources configured by remote managers and implemented by the local autonomy of a DTN node.

The DTNMA Application Management Model (AMM) [[I-D.birrane-dtn-adm](#)] defines a logical schema that can be used to represent data types and structures, autonomous controls, and other kinds of information expected to be required for the local management of a DTN node. The AMM further describes a physical data model, called the Application

Data Model (ADM), that can be defined in the context of applications to create resources in accordance with the AMM schema. These named resources can be predefined in moderated publications or custom-defined as part of the Operational Data Model (ODM) of an agent.

Every AMM resource must be uniquely identifiable. To accomplish this, an expressive naming scheme is required. The Application Resource Identifier (ARI) provides this naming scheme. This document defines the ARI, based on the structure of a Uniform Resource Identifier (URI), meeting the needs for a concise, typed, parameterized, and hierarchically organized naming convention.

1.1. Scope

The ARI scheme is based on the structure of a URI [[RFC3986](#)] in accordance with the practices outlined in [[RFC8820](#)].

ARIs are designed to support the identification requirements of the AMM logical schema. As such, this specification will discuss these requirements to the extent necessary to explain the structure and use of the ARI syntax.

This specification does not constrain the syntax or structure of any existing URI (or part thereof). As such, the ARI scheme does not impede the ownership of any other URI scheme and is therefore clear of the concerns presented in [[RFC7320](#)].

This specification does not discuss the manner in which ARIs might be generated, populated, and used by applications. The operational utility and configuration of ARIs in a system are described in other documents associated with DTN management, to include the DTNMA and AMM specifications.

This specification does not describe the way in which path prefixes associated with an ARI are standardized, moderated, or otherwise populated. Path suffixes may be specified where they do not lead to collision or ambiguity.

This specification does not describe the mechanisms for generating either standardized or custom ARIs in the context of any given application, protocol, or network.

1.2. Use of ABNF

This document defines text structure using the Augmented Backus-Naur Form (ABNF) of [[RFC5234](#)]. The entire ABNF structure can be extracted from the XML version of this document using the XPath expression:

```
'//sourcecode[@type="abnf"]'
```

The following initial fragment defines the top-level rules of this document's ABNF.

```
start = ari
```

From the document [[RFC3986](#)] the definitions are taken for pchar, unreserved, pct-encoded, path-absolute, and path-noscheme. From the document [[RFC3339](#)] the definitions are taken for date-time, full-date, and duration. From the document [[RFC5234](#)] the definitions are taken for bit, hexdig, digit, and char-val.

1.3. Use of CDDL

This document defines Concise Binary Object Representation (CBOR) structure using the Concise Data Definition Language (CDDL) of [[RFC8610](#)]. The entire CDDL structure can be extracted from the XML version of this document using the XPath expression:

```
'//sourcecode[@type="cddl"]'
```

The following initial fragment defines the top-level symbols of this document's CDDL, which includes the example CBOR content.

```
start = ari
```

```
; Limited sizes to fit the AMM data model
int32 = (int .lt 2147483648) .ge -2147483648
uint32 = uint .lt 4294967296
int64 = (int .lt 9223372036854775808) .ge -9223372036854775808
uint64 = uint .lt 18446744073709551616
```

This document does not rely on any CDDL symbol names from other documents.

1.4. 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 BCP 14 [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

Additionally, the following terms are used in this document:

Agent:

An entity being managed in the DTNMA as defined in [[I-D.ietf-dtn-dtnma](#)]. It is expected to be accessible by its Managers over a DTN.

Manager: An entity managing others in the DTNMA as defined in [[I-D.ietf-dtn-dtnma](#)]. It is expected to be accessible by its Agents over a DTN.

Application Data Model (ADM): Definitions of pre-planned objects being managed on remote agents across challenged networks. An ADM is versioned, but a single version of an ADM cannot change over time once it is registered. This is similar in function to an SMI MIB or an YANG module.

Operational Data Model (ODM): The operational configuration of an Agent, exclusive of the pre-planned objects defined by ADMs. These objects are dynamic configuration applied at runtime, either by Managers in the network or by autonomy on the Agent.

Application Resource Identifier (ARI): An identifier for any ADM or ODM managed object, as well as ad-hoc managed objects and literal values. ARIs are syntactically conformant to the Uniform Resource Identifier (URI) syntax documented in [[RFC3986](#)] and using the scheme name "ari". This is similar in function to an SMI OID or an YANG XPath expression along with parameters.

Namespace A moderated, hierarchical taxonomy of namespaces that describe a set of ADM scopes. Specifically, an individual ADM namespace is a specific sequence of ADM namespaces, from most general to most specific, that uniquely and unambiguously identify the namespace of a particular ADM.

2. ARI Purpose

ADM resources are referenced in the context of autonomous applications on an agent. The naming scheme of these resources must support certain features to inform DTNMA processing in accordance with the ADM logical schema.

This section defines the set of unique characteristics of the ARI scheme, the combination of which provides a unique utility for naming. While certain other naming schemes might incorporate certain elements, there are no such schemes that both support needed features and exclude prohibited features.

2.1. Resource Parameterization

The ADM schema allows for the parameterization of resources to both reduce the overall data volume communicated between DTN nodes and to remove the need for any round-trip data negotiation.

Parameterization reduces the communicated data volume when parameters are used as filter criteria. By associating a parameter with a data source, data characteristic, or other differentiating attribute, DTN nodes can locally process parameters to construct the minimal set of information to either process for local autonomy or report to remote managers in the network.

Parameterization eliminates the need for round-trip negotiation to identify where information is located or how it should be accessed. When parameters define the ability to perform an associative lookup of a value, the index or location of the data at a particular DTN node can be resolved locally as part of the local autonomy of the node and not communicated back to a remote manager.

2.2. Compressible Structure

The ability to encode information in very concise formats enables DTN communications in a variety of ways. Reduced message sizes increase the likelihood of message delivery, require fewer processing resources to secure, store, and forward, and require less resources to transmit.

While the encoding of an ARI is outside of the scope of this document, the structure of portions of the ARI syntax lend themselves to better compressibility. For example, DTN ADM encodings support the ability to identify resources in as few as 3 bytes by exploiting the compressible structure of the ARI.

The ARI syntax supports three design elements to aid in the creation of more concise encodings: enumerated forms of path segments, relative paths, and patterning.

2.2.1. Enumerated Path Segments

Because the ARI structure includes paths segments with stable enumerated values, each segment can be represented by either its text name or its integer enumeration. For human-readability in text form the text name is preferred, but for binary encoding and for comparisons the integer form is preferred. It is a translation done by the entity handling an ARI to switch between preferred representations (see [Section 7](#)); the data model of both forms of the ARI allows for either.

2.2.2. Relative Paths

Relative object paths and resolving are TBD. A simple method is to always use the object-containing namespace as a base URI for resolving URI References.

2.2.3. Patterning

Patterning in this context refers to the structuring of ARI information to allow for meaning data selection as a function of wildcards, regular expressions, and other expressions of a pattern.

Patterns allow for both better compression and fewer ARI representations by allowing a single ARI pattern to stand-in for a variety of actual ARIs.

This benefit is best achieved when the structure of the ARI is both expressive enough to include information that is useful to pattern match, and regular enough to understand how to create these patterns.

3. ARI Logical Structure

This section describes the components of the ARI scheme to inform the discussion of the ARI syntax in [Section 4](#). At the top-level, an ARI is one of two classes: literal or object reference. Each of these classes is defined in the following subsections.

3.1. Names, Enumerations, Comparisons, and Canonicalizations

Within the ARI logical model, there are a number of domains in which items are identified by a combination of text name and integer enumeration: ADMs, ODMs, literal types, object types, and objects. In all cases, within a single domain the text name and integer enumeration SHALL NOT be considered comparable. It is an explicit activity by any entity processing ARIs to make the translation between text name and integer enumeration (see [Section 7](#)).

Text names SHALL be restricted to begin with an alphabetic character followed by any number of other characters, as defined in the id-text ABNF symbol. This excludes a large class of characters, including non-printing characters. When represented in text form, the text name for ODMs is prefixed with a "!" character to disambiguate it from an ADM name (see [Section 3.3](#)).

For text names, comparison and uniqueness SHALL be based on case-insensitive logic. The canonical form of text names SHALL be the lower case representation.

Integer enumerations for ADMs and ODMs SHALL be restricted to a magnitude less than 2^{63} to allow them to fit within a signed 64-bit

storage. The ADM registration in [Table 6](#) reserves high-valued code points for private and experimental ADMs, while the entire domain of ODM code points (negative integers) is considered private use. Integer enumerations for literal types and object types SHALL be restricted to a magnitude less than 2^{31} to allow them to fit within a signed 32-bit storage. The registrations in [Table 3](#) and [Table 4](#) respectively Integer enumerations for objects (within an ADM or ODM) SHALL be restricted to a magnitude less than 2^{31} to allow them to fit within a signed 32-bit storage, although negative-value object enumerations are disallowed.

For integer enumerations, comparison and uniqueness SHALL be based on numeric values not on encoded forms. The canonical form of integer enumerations in text form SHALL be the shortest length decimal representation.

3.2. Literals

Literals represent a special class of ARI which are not associated with any particular ADM or ODM. A literal has no other name other than its value, but literals may be explicitly typed in order to force the receiver to handle it in a specific way.

Because literals will be based on the CBOR data model [[RFC8949](#)] and its extended diagnostic notation, a literal has an intrinsic representable data type as well as an AMM data type. The CBOR primitive types are named CDDL symbols as defined in [Section 3.3](#) of [[RFC8610](#)].

When converting from AMM literal types, the chosen CBOR type SHALL be determined by the mapping in [Table 1](#). Additionally, when handling typed literal ARIs any combination of AMM literal type and CBOR primitive type not in [Table 1](#) SHALL be considered invalid. This restriction is enforced by the CDDL defined in [Section 5](#). Additionally, when handling a literal of AMM type CBOR the well-formedness of the CBOR contained SHOULD be verified before the literal is treated as valid.

AMM Literal Type	Used CBOR Type
NULL	null
BOOL	bool
BYTE	uint
INT	int
UINT	uint
VAST	int
UVAST	uint
REAL32	float
REAL64	float

AMM Literal Type	Used CBOR Type
TEXTSTR	tstr
BYTESTR	bstr
Non-primitive types	
TP	lit-time
TD	lit-time
LABEL	lit-label
CBOR	lit-cbor
Containers	
AC	ari-collection
AM	ari-map
TBL	ari-tbl
EXECSET	exec-set
RPTSET	rpt-set

Table 1: AMM Literal Types to CBOR Types

When interpreting an untyped literal ARI, the implied AMM literal type SHALL be determined by the mapping in [Table 2](#).

CBOR Primitive Type	Implied AMM Literal Type
undefined	<i>no type</i>
null	NULL
bool	BOOL
uint	Smallest of BYTE, UINT, or UFAST to hold the value
nint	Smallest of INT, or VAST to hold the value
float16, float32	FLOAT32
float64	FLOAT64
bstr	BYTESTR
tstr	TEXTSTR

Table 2: Literal Implied and Allowed Types

3.3. Object References

Object references are composed of two parts: object identity and optional parameters. The object identity can be dereferenced to a specific object in the ADM/ODM, while the parameters provide additional information for certain types of object and only when allowed by the parameter "signature" from the ADM/ODM.

The object identity itself contains the components, described in the following subsections: namespace, object type, and object name. When encoded in text form (see [Section 4](#)), the identity components correspond to the URI path segments.

3.3.1. Namespace ID

ADM resources exist within namespaces to eliminate the possibility of a conflicting resource name, aid in the application of patterns, and improve the compressibility of the ARI. Namespaces **SHALL NOT** be used as a security mechanism to manage access. An Agent or Manager **SHALL NOT** infer security information or access control based solely on namespace information in an ARI.

Namespaces have two possible forms; one more human-friendly and one more compressible:

Text form: This form corresponds with a human-readable identifier for either an ADM or a ODM namespace. The text form is not compressible and needs to be converted to a numeric namespace based on a local registry. A text form namespace SHALL contain only URI path segment characters representing valid UTF-8 text in accordance with [[RFC3629](#)].

Numeric form: This form corresponds with a compressible value suitable for on-the-wire encoding between Manager and Agent. Sorting and matching numeric namespaces is also faster than text form. A numeric form namespaces SHALL be small enough to be represented as a 64-bit signed integer.

Independent to the form of the namespace is the issuer of the namespace, which is one of:

ADM namespace: When a namespace is associated with an ADM, its text form SHALL begin with an alphabetic character and its numeric form SHALL be a positive integer. All ADM namespaces are universally unique and, except for private or experimental use, SHOULD be registered with IANA (see [Table 6](#)).

ODM namespace: When a namespace is not associated with an ADM, its text form SHALL begin with a bang character "!" and its numeric form SHALL be a negative integer. These namespaces do not have universal registration and SHALL be considered to be private use. It is expected that runtime ODM namespaces will be allocated and managed per-user and per-mission.

3.3.2. Object Type

Due to the flat structure of an ADM, as defined in [[I-D.birrane-dtn-adm](#)], all managed objects are of a specific and unchanging type from a set of available managed object types. The preferred form for object types in text ARIs is the text name, while in binary form it is the integer enumeration (see [Section 7](#)).

The following subsection explains the form of those object identifiers.

3.3.3. Object ID

An object is any one of a number of data elements defined for the management of a given application or protocol that conforms to the ADM logical schema.

Within a single ADM or runtime namespace and a single object type, all managed objects have similar characteristics and all objects are identified by a single text name or integer enumeration. The preferred form for object names in text ARIs is the text name, while in binary form it is the integer enumeration. Any ADM-defined object will have both name and enumeration, while a runtime-defined object can have either but not both. Conversion between the two forms requires access to the original ADM, and its specific revision, in which the object was defined. A text form object name SHALL contain only URI path segment characters representing valid UTF-8 text in accordance with [[RFC3629](#)].

3.3.4. Parameters

The ADM logical schema allows many object types to be parameterized when defined in the context of an application or a protocol.

If two instances of an ADM resource have the same namespace and same object type and object name but have different parameter values, then those instances are unique and the ARIs for those instances MUST also be unique. Therefore, parameters are considered part of the ARI syntax.

The ADM logical schema defines specialized uses for the term "parameter" to disambiguate each purpose, as defined below.

Formal Parameters:

Formal parameters define the type, name, and order of the information that customizes an ARI. They represent the unchanging "signature" of the parameterized object. Because ARIs represent a *use* of an object and not its definition, formal parameters are not present in an ARI and instead are part of an object model.

Given Parameters:

Given parameters represent the data values used to distinguish different instances of a parameterized object. A given parameter is an AMM value and is represented by an ARI within the context of a parameter list (AC) or parameter map (AM). Because of necessary normalizing (of type and default value) based on formal parameters, multiple given parameters can correspond with the same meaning (see Actual Parameters below).

Additionally, there are two ways in which the value of a given parameter can be specified: parameter-by-value and parameter-by-name.

Parameter-By-Value: This method involves directly supplying the value as part of the actual parameter. It is the default method for supplying values.

Parameter-By-Name: This method involves specifying the name of an other parameter and using that other parameter's value as a substitute for the value of this parameter. This method is useful when a parameterized ARI is produced by an AMM object which itself is parameterized. The original ARI parameters contain literal ARIs with LABEL type, and when the value is produced based on input parameters the substitution is made. In this way, an actual parameter can be "flowed down" to produced values at runtime.

Actual Parameters:

Actual parameters represent a normalized set of values taken from a set of given parameters and normalized using a set of corresponding formal parameters. An actual parameter is an AMM value and is represented by an ARI.

Because normalizing can cause a given parameter to change type (in order to conform to a formal parameter type) or take on a default value (when not present in the given parameters), a single set of actual parameters can correspond with multiple options for given parameters.

4. ARI Text Form

This section defines how the data model explained in [Section 3](#) is encoded as text conforming to the URI syntax of [[RFC3986](#)]. The most straightforward text form of ARI uses an explicit scheme and an absolute path (starting with an initial slash "/"), which requires no additional context to interpret its structure.

When used within the context of a base ARI, the URI Reference form of [Section 4.4](#) can be used. In all other cases an ARI must be an absolute-path form and contain a scheme.

While this text description is normative, the ABNF schema in this section provides a more explicit and machine-parsable text schema. The scheme name of the ARI is "ari" and the scheme-specific part of the ARI follows one of the two forms corresponding to the literal-value ARI and the object-reference ARI.

```
ari = [ARIPREFIX] ari-ssp
ari-ssp = lit-ssp-struct / objref-ssp-struct

; restricted to only literals as a subset of "ari" symbol
lit-ari = [ARIPREFIX] lit-ssp-struct

ARIPREFIX = "ari:"
```

4.1. URIs and Percent Encoding

Due to the intrinsic structure of the URI, on which the text form of ARI is based, there are limitations on the syntax available to the scheme-specific-part [[RFC7595](#)]. One of these limitations is that each path segment can contain only characters in the pchar ABNF symbol defined in [[RFC3986](#)]. For most parts of the ARI this restriction is upheld by the values themselves: ADM/ODM names, literal and object type names, and object names have a limited character set all within the symbol val-seg. For literals and nested parameters though, the percent encoding of [Section 2.4](#) of [[RFC3986](#)] is needed.

The following symbol val-seg is an allowed superset of all identifiers

```
; A subset of RFC 3986 segment-nz and pchar symbols which matches all
; identifiers, primitive values, and typed-literal values
val-seg = 1*( unreserved / pct-encoded / val-delims)
val-delims = "!" / "'" / "*" / "+" / ":" / "@"

; A text name must start with an alphabetic character or underscore
id-text = (ALPHA / "_") *(ALPHA / DIGIT / "_" / "-" / ".")
; An integer enumeration must contain only digits
id-num = 1*DIGIT
```

In the ARI text examples in this document the URIs have been percent-decoded for clarity, as might be done in an ARI display and editing tool. But the actual encoded form of the human-friendly ARI `ari:"text"` is `ari:%22text%22`. Outside of literals, the safe characters which are not be percent-encoded are the structural delimiters `/()``=;`, used for parameters and ARI collections.

One other aspect of convenience for human editing of text-form ARIs is linear white space. The current ABNF pattern, staying within the URI pattern, do not allow for whitespace to separate list items or otherwise. A human editing an ARI could find it convenient to include whitespace following commas between list items, or to separate large lists across lines. Any tool that allows this kind of convenience of editing SHALL collapse any white space within a single ARI before encoding its contents.

4.2. Literals

Based on the structure of [Section 3.2](#), the text form of the literal ARI contains only a URI path with an optional AMM literal type. A literal has no concept of a namespace or context, so the path is always absolute. When the path has two segments, the first is the AMM literal type and the second is the encoded literal value. When the path has a single segment it is the encoded literal value. As a shortcut, an ARI with only a single path segment is necessarily an untyped literal so the leading slash is elided.

The text form of literal ARI has two layers of coding: the URI path structure and the type-specific value coding. These are distinguished below in the ABNF by the `lit-ssp-struct` symbol being for the general path structure and type-specific coding defined in [Section 4.2.1](#).

```
lit-ssp-struct = lit-container / lit-typeval-struct / lit-notype-struct

; More complex text (still in the "val-seg" character set) for
; non-primitive values
lit-typeval-struct = "/" lit-type "/" val-seg
; Type is restricted to valid AMM literal types
lit-type = id-text / id-num

; These use different symbols than lit-typeval-struct because of recursi
lit-container = lit-ac / lit-am / lit-tbl / lit-execset / lit-rptset
lit-ac = "/" ("AC" / "17") "/" ari-collection
lit-am = "/" ("AM" / "18") "/" ari-map
lit-tbl = "/" ("TBL" / "19") "/" ari-tbl
lit-execset = "/" ("EXECSET" / "20") "/" exec-set
lit-rptset = "/" ("RPTSET" / "21") "/" rpt-set

; The untyped value is a subset of the "lit-notype" symbol
lit-notype-struct = val-seg
```

4.2.1. Typed Literal Values

The definition in this section is a specialization of the `lit-typeval-struct` structure for specific literal types.

An ARI encoder or decoder SHALL handle both text name and integer enumeration forms of the `lit-type` symbol. When present and able to be looked up, the literal type SHOULD be a text name.

The text form of typed values SHALL be one of the following, based on the associated `lit-type` value in the ARI:

NULL:

This type contains only the single the fixed value "null".

```
null = "null"
```

BOOL:

This type contains the two fixed values "true" and "false".

```
bool = "true" / "false"
```

BYTE, INT, UINT, VAST, or UVAST:

The integer types, signed or unsigned, match the following integer ABNF symbol:

```
integer = int-dec / int-bin / int-hex
int-dec = optsign 1*DIGIT
int-bin = optsign "0b" 1*BIT
int-hex = optsign "0x" 1*(HEXDIG HEXDIG)
optsign = ["+" / "-"]
```

REAL32 or REAL64:

The floating-point types match the following float ABNF symbol, which includes a specialized text form that avoids decimal or exponential conversion of the underlying value. The raw floating point text form SHALL be the prefix "0fx" followed by the hexadecimal encoding of the network-byte-order binary encoding in accordance with [[IEEE.754-2019](#)] for binary16, binary32, or binary64 values.

```
float = float-exp / float-dec / float-raw
float-exp = optsign 1*DIGIT ["." 1*DIGIT] "e" optsign 1*DIGIT
float-dec = optsign 1*DIGIT "." 1*DIGIT
; Custom encoding of IEEE-754 binary content
; The hex data will really only be 2, 4, or 8 bytes
float-raw = optsign "0fx" *(HEXDIG HEXDIG)
```

TEXTSTR:

The text string type matches the following tstr ABNF symbol after percent-decoding, which is consistent with the CBOR Diagnostic Notation definition. As an alternative, if the text value matches the id-text character set it can be encoded directly without quoting.

```
; double-quoted and escaped text
tstr = tstr-quoted / id-text
tstr-quoted = DQUOTE *(%x20-21 / %x23-7E) DQUOTE
```

BYTESTR:

The byte string type matches the following bstr ABNF symbol after percent-decoding, which is consistent with the CBOR Diagnostic Notation definition. All non-raw byte strings are encoded in accordance with [[RFC4648](#)].


```
bstr = bstr-raw / bstr-b16 / bstr-b64
bstr-raw = SQUOTE *(%x20-21 / %x23-7E) SQUOTE
bstr-b16 = "h" SQUOTE 1*(HEXDIG HEXDIG) SQUOTE
bstr-b32 = "b32" SQUOTE 1*(ALPHA / %x32-37) *EQ SQUOTE
bstr-b64 = "b64" SQUOTE 1*(ALPHA / DIGIT / "+" / "/") *EQ SQUOTE
SQUOTE = "'"
EQ = "="
```

TP:

This type uses either the date-time ABNF symbol of [Appendix A](#) of [\[RFC3339\]](#) always in the "Z" time-offset, or as a numeric representation of the relative time from the DTN epoch. This text is unquoted and, to avoid percent encoding, this text form MAY omit the separator characters "-" and ":".

```
tp-val = date-time / float / integer
```

TD:

This type uses either the duration ABNF symbol of [Appendix A](#) of [\[RFC3339\]](#) with a positive or negative sign prefix, or as a numeric representation of the relative time value. This text is unquoted and due to the constraints on the value need not be percent encoded.

```
td-val = duration / float / integer
```

LABEL:

This type uses the id-text ABNF symbol from this document. This text is unquoted and due to the constraints on the value need not be percent encoded.

CBOR:

This type uses the following emb-cbor ABNF symbol for its value. The first encoding choice uses the percent encoded form of CBOR extended diagnostic notation of [Appendix G](#) of [\[RFC8610\]](#). The second encoding is the raw byte string of the embedded CBOR.

```
emb-cbor = cbor-diag / bstr
cbor-diag = "<<" *(%x20-3B / %x3D / %x3F-7E) ">>"
```

AC:

This type uses the following ari-collection ABNF symbol for its value. Each item of the collection is an already-percent-encoded text-form ARI.

```
; A comma-separated list of any form of ARI with enclosure
ari-collection = "(" ari *("," ari) ")"
```

AM:

This type uses the following ari-map ABNF symbol for its value. Each key of the map is an already-percent-encoded text-form literal-value ARI. Each value of the map is an already-percent-encoded text-form ARI.

```
; A comma-separated list of pairs, each delimited by equal-sign
; with literal-only keys
ari-map = "(" ari-map-pair *("," ari-map-pair) ")"
ari-map-pair = lit-notype "=" ari
```

TBL:

This type uses the following ari-tbl ABNF symbol for its value. The value is prefixed by the number of columns in the table, followed by an AC representing each separate row in the table. Each item of the table is an already-percent-encoded text-form ARI.

```
ari-tbl = "c=" *DIGIT ";" *ari-collection
```

Note that although the TBL uses the same syntax as AC for encoding each row, the value itself is not necessarily internally represented by a sequence of AC values.

EXECSET:

This type uses the following exec-set ABNF symbol for its value. The value is prefixed by the optional nonce value for the associated execution(s), followed by an AC representing each item to be executed (either a CTRL reference, MAC-producing reference, or MAC value). Each item of the targets is an already-percent-encoded text-form ARI.

```
exec-set = "n=" exec-nonce ";" exec-targets
exec-nonce = null / integer / bstr
exec-targets = ari-collection
```

Note that although the EXECSET uses the same syntax as AC for encoding each row, the value itself is not necessarily internally represented by an AC.

RPTSET:

This type uses the following rpt-set ABNF symbol for its value. The value consists of the optional nonce value for the associated execution(s), a reference absolute time for all contained reports, and the list of contained reports. Each contained report consists of a relative creation time for the report (relative to the RPTSET reference time), a reference to the source of the report, and an

AC representing each item in the report. Each item of a report is an already-percent-encoded text-form ARI, as is the source reference. The nonce and timestamps are untyped values.

```
rpt-set = "n=" exec-nonce ";r=" tp-val ";" *rpt-container
rpt-container = "(t=" td-val ";s=" ari ";" rpt-items ")"
rpt-items = ari-collection
```

Note that although the report container uses the same syntax as AC for encoding each row, the items itself is not necessarily internally represented by an AC.

Some examples of typed literal values are below. The represented values for TP, TD, and CBOR types are the same just with different text representations.

```
ari:/BOOL/true
ari:/UINT/10
ari:/VAST/10
ari:/LABEL/name
ari:/TP/20230102T030405Z
ari:/TP/2023-01-02T03:04:05Z
ari:/TP/725943845
ari:/TD/+PT1H
ari:/TD/3600
ari:/CBOR/h'0a'
ari:/CBOR/%3C%3C10%3E%3E
```

These are typed container values with text form for AC and AM lists, here not percent-encoded:

```
ari:/AC/(1,2,3)
ari:/AM/(1=2,2=4,3=9)
ari:/TBL/c=3;(1,true,%22A%22)(2,false,%22B%22)
```

4.2.2. Untyped Literal Values

The definition in this section is a specialization of the lit-notype-struct structure for specific primitive types.

When untyped, the literal value SHALL be one of the primitive types named by the lit-notype ABNF symbol below. The separate value encodings use symbols from [Section 4.2.1](#).

The order of this symbol sequence is significant because the unquoted tstr must be matched after the enumerated types of undefined, null, and bool.

```
lit-notype = undefined / null / bool / float / integer / tstr / bstr
; the undefined value is only valid as an untyped literal
undefined = "undefined"
```

Some example of the forms for a literals are below. These first are untyped literal values:

```
ari:undefined
ari:true
ari:10
ari:%27bytes%27
ari:h%276869%27
ari:%22text%22
```

4.2.3. Preferred Encodings

Several of the literal types defined in [Section 4.2.1](#) allow the same AMM value to be represented by multiple logically equivalent encodings. Because these encodings are not equivalent in represented text size or processing needs it is useful for an ARI processor to allow a user to determine preferred encodings when processing text-form ARIs.

Integer values: These options correspond to the BYTE, INT, UINT, VAST, and UVAST types and untyped int values. These values can be encoded as either base-2, base-10, or base-16. In uses which are more resource constrained and less human-facing a processor **MAY** encode integers only in base-16.

Floating point values: These options correspond to the REAL32 and REAL64 types and untyped float values. These values can be encoded as either dotted-decimal, exponential, or hex-encoded raw bytes. In uses which are more resource constrained and less human-facing a processor **MAY** encode floats only as raw bytes.

Byte string values: These options correspond to the BYTESTR type and untyped bstr values. These values can be encoded as either raw bytes, base-16, base-32, or base-64. In uses which are more resource constrained and less human-facing a processor **MAY** encode byte strings only as base-16.

Time values: These options correspond to the TP and TD types. These values can be encoded as either delimited text, non-delimited text, or decimal numbers. In uses which are more resource constrained and less human-facing a processor **MAY** encode time values only as decimal numbers.

Embedded CBOR values: These options correspond to the CBOR type. These values can be encoded as either CBOR Diagnostic Notation or as a byte string. In uses which are more resource constrained and

less human-facing a processor **MAY** encode CBOR values only as byte strings.

4.3. Object References

Based on the structure of [Section 3.3](#), the text form of the object reference ARI contains a URI with three path segments corresponding to the namespace-id, object-type, and object-id. Those three segments (excluding parameters as defined below) are referred to as the object identity.

An ARI encoder or decoder SHALL handle both text name and integer enumeration forms of the namespace-id, object-type, and object-id.

The final segment containing the object-id MAY contain parameters enclosed by parentheses "(" and ")". There is no semantic distinction between the absence of parameters and the empty parameter list. The contents of the parameters SHALL be interpreted as a literal AC or AM in accordance with [Section 4.2](#).

The parameters as a whole SHALL be the percent encoded form of the constituent ARIs, excluding the structural delimiters /(),=.

Implementations are advised to be careful about the percent encoded vs. decoded cases of each of the nested ARIs within parameters to avoid duplicate encoding or decoding. It is recommended to dissect the parameters and ARI collections in their encoded form first, and then to dissect and percent decode each separately and recursively.

```
ari:/adm-a/EDD/someobj
ari:/adm-a/CTRL/otherobj(true,3)
ari:/adm-a/CTRL/otherobj(%22a%20param%22,/UINT/10)
ari:/41/-1/0
```

The object-reference ARI has a corresponding ABNF definition of:

```
objref-ssp-struct = (obj-ident / obj-ident) [paramlist]
; The object identity can be used separately than parameters
obj-ident = ("/" ns-id / ".") "/" obj-type "/" obj-id

; Parameters are either AC or AM
paramlist = ari-collection / ari-map

ns-id = ns-adm / ns-odm
ns-adm = id-text ["@" full-date] / id-num
ns-odm = ("!" id-text) / ("- " id-num)
; Type is restricted to valid AMM object types
obj-type = id-text / ("- " id-num)
obj-id = id-text / id-num
```

4.4. URI References

The text form of ARI can contain a URI Reference, as defined in [Section 3](#) of [[RFC3986](#)], which can only be resolved using a base URI using the algorithm defined in [Section 5](#) of [[RFC3986](#)]. When resolving nested ARI content, the base URI of any interior resolution is the next-outer ARI in the nested structure. The outermost ARI SHALL NOT be a URI Reference because it will have no base URI to resolve with.

Because a relative-path ARI with no path separators is considered to be an untyped literal, an ARI reference SHALL contain at least one path separator. For the case where the ARI reference is to a sibling object from the base URI the relative path SHOULD be of the form "./" to include the path separator.

When resolving nested ARI content, the parameters of the URI reference SHALL be preserved in the resolved ARI. This behavior is equivalent to the query parameter portion when resolving a generic URI reference.

```
; Relative ARI must be resolved before interpreting
relative-ari = path-nonempty [paramlist]
```

```
; Non-empty absolute or relative path
path-nonempty = path-absolute / path-noscheme
```

5. ARI Binary Form

This section defines how the data model explained in [Section 3](#) is encoded as a binary sequence conforming to the CBOR syntax of [[RFC8949](#)]. Within this section the term "item" is used to mean the CBOR-decoded data item which follows the logical model of CDDL [[RFC8610](#)].

The binary form of the URI is intended to be used for machine-to-machine interchange so it is missing some of the human-friendly shortcut features of the ARI text form from [Section 4](#). It still follows the same logical data model so it has a one-for-one representation of all of the styles of text-form ARI.

A new CBOR tag TBD999999 has been registered to indicate that an outer CBOR item is a binary-form ARI. This is similar in both syntax and semantics to the "ari" URI scheme in that for a nested ARI structure, only the outer-most ARI need be tagged. The inner ARIs are necessarily interpreted as such based on the nested ARI schema of this section.

While this text description is normative, the CDDL schema in this section provides a more explicit and machine-parsable binary schema.

```
; An ARI can be tagged if helpful
ari = ari-notag / #6.999999(ari-notag)
ari-notag = lit-ari / objref-ari
```

5.1. Intermediate CBOR

The CBOR item form is used as an intermediate encoding between the ARI data and the ultimate binary encoding. When decoding a binary form ARI, the CBOR must be both "well-formed" according to [\[RFC8949\]](#) and "valid" according to the CDDL model of this specification. Implementations are encouraged, but not required, to use a streaming form of CBOR encoder/decoder to reduce memory consumption of an ARI handler. For simple implementations or diagnostic purposes, a two stage conversion between ARI--CBOR and CBOR--binary can be more easily understood and tested.

5.2. Literals

Based on the structure of [Section 3.2](#), the binary form of the literal ARI contains a data item along with an optional AMM literal type identifier. In order to keep the encoding as short as possible, the untyped literal is encoded as the simple value itself. Because the typed literal and the object-reference forms uses CBOR array framing, this framing is used to disambiguate from the pure-value encoding of the lit-notype CDDL symbol.

When present, the literal type SHALL be an integer enumeration. When typed, the decoded literal value SHALL be a valid CBOR item conforming to the AMM literal type definition of the \$lit-typeval CDDL socket. When untyped, the decoded literal value SHALL be one of the primitive types named by the lit-notype CDDL symbol.

```

lit-ari = lit-typeval / lit-notype

; undefined value is only allowed as non-typed literal
lit-notype = undefined / null / bool / int / float / tstr / bstr

lit-typeval = $lit-typeval .within lit-typeval-struct
lit-typeval-struct = [
  lit-type: lit-type-id,
  lit-value: any
]
lit-type-id = (int32 .ge 0)

; IANA-assigned literal types
$lit-typeval /= [0, null]
$lit-typeval /= [1, bool]
$lit-typeval /= [2, uint .size 1] ; 1-byte
$lit-typeval /= [4, int32] ; 4-byte
$lit-typeval /= [5, uint32] ; 4-byte
$lit-typeval /= [6, int64] ; 8-byte
$lit-typeval /= [7, uint64] ; 8-byte
$lit-typeval /= [8, float16 / float32]
$lit-typeval /= [9, float16 / float32 / float64]
$lit-typeval /= [10, tstr]
$lit-typeval /= [11, bstr]

; Absolute timestamp as seconds from DTN Time epoch
$lit-typeval /= [12, lit-time]
; Relative time interval as seconds
$lit-typeval /= [13, lit-time]
lit-time = int / time-fraction
; Same structure as tag #4 "decimal fraction" but limited in domain
time-fraction = [
  exp: (-9 .. 9) .within int,
  mantissa: int,
]

; Parameter label
$lit-typeval /= [14, lit-label]
lit-label = tstr .regex "[A-Za-z].*"

; Embedded CBOR item
$lit-typeval /= [15, lit-cbor]
lit-cbor = bstr .cbor any

; Literal type ID value
$lit-typeval /= [16, lit-type-id]

; Ordered list of ARIs
$lit-typeval /= [17, ari-collection]

```



```

ari-collection = [*ari-notag]

; Map from untyped literal to ARI
$lit-typeval /= [18, ari-map]
ari-map-key = lit-notype
ari-map = {*ari-map-key => ari-notag}

; Table of ARIs
$lit-typeval /= [19, ari-tbl]
ari-tbl = [
  ncol: uint, ; Number of columns in the table
  *tbl-row ; All rows are concatenated in the single array
]
tbl-row = (*ari-notag)

; Execution-Set
$lit-typeval /= [19, exec-set]
exec-set = [
  nonce,
  exec-targets,
]
nonce = bstr / uint / null
exec-targets = (*objref-ari)

; Reporting-Set
$lit-typeval /= [20, rpt-set]
rpt-set = [
  nonce,
  ref-time: lit-time, ; Interpreted as a TP absolute time
  *rpt-container
]
rpt-container = [
  rel-time: lit-time ; Interpreted as a TD relative to ref-time
  source: objref-ari,
  rpt-items
]
rpt-items = (*ari-notag)

```

Some example of the forms for a literal are below. These first are untyped primitive values:

true

"text"

10

And these are typed values:

[4, 10]

[15, <<10>>]

The literal-value ARI has a corresponding CDDL definition of:

5.3. Object References

Based on the structure of [Section 3.3](#), the binary form of the object reference ARI is a CBOR-encoded item. An ARI SHALL be encoded as a CBOR array with at least three items corresponding to the namespace-id, object-type, and object-id. Those three items are referred to as the object identity. The optional fourth item of the array is the parameter list.

The namespace-id SHALL be present only as an integer enumeration. The object-type SHALL be present only as an integer enumeration. The object-id SHALL be present as either a text name or an integer enumeration. The processing of text name object identity components by an Agent is optional and SHALL be communicated to any associated Manager prior to encoding any ARIs for that Agent.

When present, the parameters SHALL be either the ari-collection or ari-map structure. In other words, just the value-portion of the AC or AM typed literal because no other disambiguation needs to be made for the parameter type.

An example object reference without parameters is:

[41, -1, 0]

Another example object reference with parameters is:

[41, -2, 3, ["a param", [4, 10]]]

The object-reference ARI has a corresponding CDDL definition of:

```

; Type-agnostic structure of object-reference
objref-ari = $objref-ari .within objref-ari-struct
objref-ari-struct = [
    obj-ident<obj-type-id>,
    ?params: ari-collection / ari-map
]

; Identity of a single object
obj-ident<obj-type> = (
    ns-id,
    obj-type,
    obj-id,
)
; The null-value namespace means relative to a context ADM/ODM
ns-id = int64 / null
obj-type-id = (int32 .lt 0)
obj-id = (int32 .ge 0) / tstr

; Generic usable for restricting objref-ari by type
objref-type<obj-type> = [
    obj-ident<obj-type>,
    ?params: ari-collection / ari-map
]

; IANA-assigned object types
CONST = -2
$objref-ari /= objref-type<CONST>
CTRL = -3
$objref-ari /= objref-type<CTRL>
EDD = -4
$objref-ari /= objref-type<EDD>
OPER = -6
$objref-ari /= objref-type<OPER>
SBR = -8
$objref-ari /= objref-type<SBR>
TBR = -10
$objref-ari /= objref-type<TBR>
VAR = -11
$objref-ari /= objref-type<VAR>
TYPEDEF = -12
$objref-ari /= objref-type<TYPEDEF>

```

5.4. URI References

The binary form of ARI does not provide the ability to encode a URI Reference. When a text-form ARI contains a relative path, it must be fully resolved before it can be translated to a binary-form ARI.

6. ARI Patterns

Because the [ARI logical structure \(Section 3\)](#) uses path segments to delimit the components of the absolute path, and due to the restrictions of the ARI path segment content, it is possible for URI reserved characters to be able to provide wildcard-type patterns.

Although the form is similar, an ARI Pattern is not itself an ARI and they cannot be used interchangeably. The context used to interpret and match an ARI Pattern SHALL be explicit and separate from that used to interpret and dereference an ARI.

While an ARI is used to dereference to a specific managed object and invoke behavior on that object, an ARI Pattern is used solely to perform matching logic against text-form ARIs. The ARI Pattern SHALL NOT ever take the form of a URI Reference; only as an absolute URI. An ARI Pattern SHALL NOT ever contain parameters, only identity.

An ARI Pattern has no optional path segments. When used as a literal ARI pattern the path SHALL have two segments. When used as an object-reference ARI pattern the path SHALL have three segments.

The single-wildcard is the only defined segment pattern and a segment can either be a real ID or a single wildcard. Because an ARI Pattern is just used to match text-form ARIs it has no specific restrictions on enumerated segment text the way a valid ARI does.

The ABNF for the ARI Pattern is given here:

```
ari-pat = "ari:" ari-pat-ssp
ari-pat-ssp = ari-pat-literal / ari-pat-objref

ari-pat-literal = "/" id-pat "/" id-pat
ari-pat-objref = "/" id-pat "/" id-pat "/" id-pat

; The non-wildcard symbol is the same as ARI syntax
id-pat = wildcard / (*pchar)
wildcard = "*"

```

6.1. ARI Matching

An ARI Pattern SHALL be considered to match an ARI when, after removing the ARI's parameters, each component of the ARI Pattern matches the corresponding ARI path component. Each pattern component SHALL be considered to match according to the following rules:

Specific value: The pattern component SHALL be compared with the ARI component after both are percent-decoded in accordance with [Section 2.1](#) of [[RFC3986](#)] and UTF-8 decoded in accordance with [[RFC3629](#)].

Single-segment wildcard:

The pattern component SHALL be considered to match with any ARI component, if present, including an empty component.

7. Transcoding Considerations

When translating literal types into text form and code point lookup tables are available, the literal type SHOULD be converted to its text name. When translating literal types from text form and code point lookup tables are available, the literal type SHOULD be converted from its text name. The conversion between AMM literal type name and enumeration requires a lookup table based on the registrations in [Table 3](#).

When translating literal values into text form, it is necessary to canonicalize the CBOR extended diagnostic notation of the item. The following applies to generating text form from CBOR items:

- *The canonical text form of CBOR bool values SHALL be the forms identified in [Section 8](#) of [\[RFC8949\]](#).
- *The canonical text form of CBOR int and float values SHALL be the decimal form defined in [Section 8](#) of [\[RFC8949\]](#).
- *The canonical text form of CBOR tstr values SHALL be the definite-length, non-concatenated form defined in [Section 8](#) of [\[RFC8949\]](#).
- *The canonical text form of CBOR bstr values SHALL be the definite-length, base16 ("h" prefix), non-concatenated form defined in [Section 8](#) of [\[RFC8949\]](#).
- *When presenting the AMM literal type of CBOR the values SHALL be the embedded CBOR form defined in [Appendix G.3](#) of [\[RFC8610\]](#).

When translating object references into text form and code point lookup tables are available, any enumerated item SHOULD be converted to its text name. When translating object references from text form and code point lookup tables are available, any enumerated item SHOULD be converted from its text name. The conversion between AMM object-type name and enumeration requires a lookup table based on the registrations in [Table 4](#). The conversion between name and enumeration for either namespace-id or object-id require lookup tables based on ADMs and ODMs known to the processing entity.

8. Interoperability Considerations

DTN challenged networks might interface with better resourced networks that are managed using non-DTN management protocols. When this occurs, the federated network architecture might need to define

management gateways that translate between DTN and non-DTN management approaches.

NOTE: It is also possible for DTN management be used end-to-end because this approach can also operate in less challenged networks. The opposite is not true; non-DTN management approaches should not be assumed to work in DTN challenged networks.

Where possible, ARIs should be translatable to other, non-DTN management naming schemes. This translation might not be 1-1, as the features of the ADM may differ from features in other management naming schemes. Therefore, it is unlikely that a single naming scheme can be used for both DTN and non-DTN management.

9. Security Considerations

Because ADM and ODM namespaces are defined by any entity, no security or permission meaning can be inferred simply from the expression of namespace.

10. IANA Considerations

This section provides guidance to the Internet Assigned Numbers Authority (IANA) regarding registration of schema and namespaces related to ARIs, in accordance with BCP 26 [[RFC1155](#)].

10.1. URI Schemes Registry

This document defines a new URI scheme "ari" in [Section 4](#). A new entry has been added to the "URI Schemes" registry [[IANA-URI](#)] with the following parameters.

Scheme name:

ari

Status:

Permanent

Applications/protocols that use this scheme name:

The scheme is used by DTNMA Managers and Agents to identify managed objects.

Contact:

IETF Chair <chair@ietf.org>

Change controller:

IESG <iesg@ietf.org>

Reference:

[Section 4](#) of [This document].

10.2. CBOR Tags Registry

This document defines a new CBOR tag TBD999999 in [Section 5](#). A new entry has been added to the "CBOR Tags" registry [[IANA-CBOR](#)] with the following parameters.

Tag:

TBD999999

Data Item:

multiple

Semantics:

Used to tag a binary-form DTNMA ARI

Reference:

[Section 5](#) of [This document].

10.3. DTN Management Architecture Parameters

This document defines several new sub-registries within a new "DTN Management Architecture Parameters" registry.

This document defines a new sub-registry "Literal Types" within the "DTN Management Architecture Parameters" registry [[IANA-DTNMA](#)] containing initial entries from [Table 3](#). Enumerations in this sub-registry SHALL be non-negative integers representable as CBOR uint type with an argument shorter than 4-bytes. Names in this sub-registry SHALL be unique among all entries in this and the "Managed Object Types" sub-registry. The registration procedure for this sub-registry is Specification Required.

Enumeration	Name	Description	Reference
255	LITERAL	A reserved type name for the union of all possible literal types.	[This document]
0	NULL	The singleton null value.	[This document]
1	BOOL	A native boolean true or false value.	[This document]
2	BYTE	An 8-bit unsigned integer.	[This document]
4	INT	A 32-bit signed integer.	[This document]
5	UINT	A 32-bit unsigned integer.	[This document]
6	VAST	A 64-bit signed integer.	[This document]
7	UFAST	A 64-bit unsigned integer.	

Enumeration	Name	Description	Reference
			[This document]
8	REAL32	A 32-bit [IEEE.754-2019] floating point number.	[This document]
9	REAL64	A 64-bit [IEEE.754-2019] floating point number.	[This document]
10	TEXTSTR	A text string composed of (unicode) characters.	[This document]
11	BYTESTR	A byte string composed of 8-bit values.	[This document]
12	TP	An absolute time point (TP).	[This document]
13	TD	A relative time difference (TD) with a sign.	[This document]
14	LABEL	A text label of a parent object parameter. This is only valid in a nested parameterized ARI.	[This document]
15	CBOR	A byte string containing an encoded CBOR item. The structure is opaque to the Agent but guaranteed well-formed for the ADM using it.	[This document]
16	ARITYPE	An integer value representing one of the code points in this Literal Types table or the Object Types table.	[This document]
17	AC	An array containing an ordered list of ARIs.	[This document]
18	AM	A map containing keys of primitive ARIs and values of ARIs.	[This document]
19	TBL	A two-dimensional table containing cells of ARIs.	[This document]
20	EXECSET	A structure containing values to be executed by an Agent.	[This document]
21	RPTSET	A structure containing reports of values sampled from an Agent.	[This document]
22 to 254, 256 to 65279		<i>Unassigned</i>	
65280 to 2147483647		Enumerations that are 2^{16} - 2^8 and larger are reserved for private or experimental use.	[This document]

Table 3: Literal Types

This document defines a new sub-registry "Managed Object Types" within the "DTN Management Architecture Parameters" registry [IANA-DTNMA] containing initial entries from Table 4. Enumerations in this sub-registry SHALL be negative integers representable as CBOR

nint type with an argument shorter than 4-bytes. Names in this sub-registry SHALL be unique among all entries in this and the "Literal Types" sub-registry. The registration procedure for this sub-registry is Specification Required.

Enumeration	Name	Description	Reference
-256	OBJECT	A reserved type name for the union of all possible object types.	[This document]
-2	CONST	Constant	[This document]
-3	CTRL	Control	[This document]
-4	EDD	Externally Defined Data	[This document]
-6	OPER	Operator	[This document]
-8	SBR	State-Based Rule	[This document]
-10	TBR	Time-Based Rule	[This document]
-11	VAR	Variable	[This document]
-12	TYPEDEF	ADM-defined type	[This document]
-13 to -255, -257 to -65280		<i>Unassigned</i>	
-65281 to -2147483648		Enumerations that are -1-(2**16-2**8) and larger are reserved for private or experimental use.	[This document]

Table 4: Managed Object Types

This document defines a new sub-registry "Application Data Models" within the "DTN Management Architecture Parameters" registry [[IANA-DTNMA](#)] containing initial entries from [Table 6](#). Enumerations in this sub-registry are non-negative integers representable as CBOR uint type with an argument shorter than 8-bytes. The registration procedures for this sub-registry are indicated in [Table 5](#).

Enumeration Range	Registration Procedure
1 to 24	Standards Action With Expert Review
25 to 255	Private Use
256 to 4294967296	Specification Required

Table 5: Application Data Models Registration Procedures

Enumeration	Name	Reference	Notes
0		[This document]	Value zero is reserved.
1 to 24			<i>Unassigned</i>
25 to 255		[This document]	Enumerations with 8-bit size are reserved for Private Use
256 to 4294967296			<i>Unassigned</i>
4294967296 and larger		[This document]	Enumerations that are larger than 32-bit are reserved for private or experimental use.

Table 6: Application Data Models

11. References

11.1. Normative References

- [IANA-CBOR] IANA, "Concise Binary Object Representation (CBOR) Tags", <<https://www.iana.org/assignments/cbor-tags/>>.
- [IANA-DTNMA] IANA, "Delay-Tolerant Networking Management Architecture (DTNMA) Parameters", <<https://www.iana.org/assignments/TBD/>>.
- [IANA-URI] IANA, "Uniform Resource Identifier (URI) Schemes", <<https://www.iana.org/assignments/uri-schemes/>>.
- [IEEE.754-2019] IEEE, "IEEE Standard for Floating-Point Arithmetic", IEEE IEEE 754-2019, DOI 10.1109/IEEESTD.2019.8766229, 18 July 2019, <<https://ieeexplore.ieee.org/document/8766229>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3339] Klyne, G. and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, DOI 10.17487/RFC3339, July 2002, <<https://www.rfc-editor.org/info/rfc3339>>.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <<https://www.rfc-editor.org/info/rfc3629>>.
- [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, <<https://www.rfc-editor.org/info/rfc3986>>.

- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <<https://www.rfc-editor.org/info/rfc4648>>.
- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, DOI 10.17487/RFC5234, January 2008, <<https://www.rfc-editor.org/info/rfc5234>>.
- [RFC7595] Thaler, D., Ed., Hansen, T., and T. Hardie, "Guidelines and Registration Procedures for URI Schemes", BCP 35, RFC 7595, DOI 10.17487/RFC7595, June 2015, <<https://www.rfc-editor.org/info/rfc7595>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8949] Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", STD 94, RFC 8949, DOI 10.17487/RFC8949, December 2020, <<https://www.rfc-editor.org/info/rfc8949>>.
- [RFC9171] Burleigh, S., Fall, K., and E. Birrane, III, "Bundle Protocol Version 7", RFC 9171, DOI 10.17487/RFC9171, January 2022, <<https://www.rfc-editor.org/info/rfc9171>>.

11.2. Informative References

- [RFC1155] Rose, M. and K. McCloghrie, "Structure and identification of management information for TCP/IP-based internets", STD 16, RFC 1155, DOI 10.17487/RFC1155, May 1990, <<https://www.rfc-editor.org/info/rfc1155>>.
- [RFC4838] Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., and H. Weiss, "Delay-Tolerant Networking Architecture", RFC 4838, DOI 10.17487/RFC4838, April 2007, <<https://www.rfc-editor.org/info/rfc4838>>.
- [RFC7320] Nottingham, M., "URI Design and Ownership", RFC 7320, DOI 10.17487/RFC7320, July 2014, <<https://www.rfc-editor.org/info/rfc7320>>.
- [RFC8610] Birkholz, H., Vigano, C., and C. Bormann, "Concise Data Definition Language (CDDL): A Notational Convention to Express Concise Binary Object Representation (CBOR) and JSON Data Structures", RFC 8610, DOI 10.17487/RFC8610, June 2019, <<https://www.rfc-editor.org/info/rfc8610>>.

[RFC8820]

Nottingham, M., "URI Design and Ownership", BCP 190, RFC 8820, DOI 10.17487/RFC8820, June 2020, <<https://www.rfc-editor.org/info/rfc8820>>.

[I-D.ietf-dtn-dtnma] Birrane, E. J., Heiner, S., and E. Annis, "DTN Management Architecture", Work in Progress, Internet-Draft, draft-ietf-dtn-dtnma-08, 10 December 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-dtn-dtnma-08>>.

[I-D.birrane-dtn-adm] III, E. J. B., Sipos, B., and J. Ethier, "DTNMA Application Management Model (AMM) and Data Models", Work in Progress, Internet-Draft, draft-birrane-dtn-adm-06, 3 January 2024, <<https://datatracker.ietf.org/api/v1/doc/document/draft-birrane-dtn-adm/>>.

Appendix A. Examples

The examples in this section rely on the ADM and ODM definitions in [Table 7](#) and [Table 8](#) respectively.

Enumeration	Name
10	adm10
20	adm20

Table 7: Example ADMs

Enumeration	Name
-10	odm10

Table 8: Example ODMs

Given those namespaces, the example types are listed in [Table 10](#) and objects are listed in [Table 9](#) where the Namespace column uses the ARI text form convention.

Namespace	Object Type	Enumeration	Name	Signature
adm10	EDD	3	num_bytes	()
adm10	CTRL	2	do_thing	(AC targets, UINT count)
adm10	CONST	1	rpt_with_param	(ARI var, TEXTSTR text)
!odm10	VAR	1	my_counter	()

Table 9: Example Objects

Namespace	Type Name	Enumeration	Summary
adm10	distance	1	A specialization of uint with scale of 1.0 and unit of meter.

Table 10: Example ADM Types

Each of the following examples illustrate the comparison of ARI forms in different situations, covering the gamut of what can be expressed by an ARI.

A.1. Primitive-Typed Literal

This is the literal value 4 interpreted as a 32-bit unsigned integer. The ARI text (which is identical to its percent-encoded form) is:

```
ari:/UINT/4
```

which is translated to enumerated form:

```
ari:/5/4
```

and converted to CBOR item:

```
[5, 4]
```

and finally to the binary string of:

```
0x820504
```

A.2. Timestamp Literal

This is the timestamp "2000-01-01T00:16:40Z" which is DTN Time epoch plus 1000 seconds. The ARI text (which is identical to its percent-encoded form) is:

```
ari:/TP/20000101T001640Z
```

which is translated to enumerated form:

```
ari:/12/1000
```

and converted to CBOR item:

```
[12, 1000]
```

and finally to the binary string of:

```
0x820c1a000f4240
```

A.3. Semantic-Typed Literal

This is the literal value 20 interpreted as a semantic type distance from adm10. The ARI text (which is identical to its percent-encoded form) is:

```
ari:/adm10/TYPEDEF/distance(20)
```

which is translated to enumerated form:

```
ari:/10/-12/1(20)
```

and converted to CBOR item:

```
[10, -12, 1, [20]]
```

and finally to the binary string of:

```
0x840a2b018114
```

A.4. Complex CBOR Literal

This is a literal value embedding a complex CBOR structure. The CBOR diagnostic expression being encoded is

```
<<{"test": [3, 4.5]}>>
```

which can be directly percent encoded as

```
ari:/CBOR/%3C%3C%7B%22test%22%3A%5B3%2C4.5%5D%7D%3E%3E
```

The embedded item can further be CBOR-encoded to a byte string and percent-encoded, along with a translated type enumeration of:

```
ari:/15/h%27A164746573748203F94480%27
```

and converted to CBOR item (note the byte string is no longer text-encoded):

```
[15, h'A164746573748203F94480']
```

and finally to the binary string of:

```
0x820F4BA164746573748203F94480
```

A.5. Non-parameterized Object Reference

This is a non-parameterized num_bytes object in the ADM namespace. The ARI text (which is identical to its percent-encoded form) is:

```
ari:/adm10/edd/num_bytes
```

which is translated to enumerated form:

```
ari:/10/-4/3
```

and converted to CBOR item:

```
[10, -4, 3]
```

and finally to the binary string of:

```
0x830A2303
```

A.6. Parameterized Object Reference

This is an parameterized do_thing object in the ADM namespace. Additionally, the parameters include two relative-path ARI References to other objects in the same ADM, which are resolved after text-decoding. The ARI text (which is identical to its percent-encoded form) is:

```
ari:/adm10/ctrl/do_thing(/AC/(./edd/num_bytes,/!odm10/var/my_counter),3)
```

which is translated to enumerated and resolved form:

```
ari:/10/-3/2(/17(/10/-4/3,-10/-11/1),3)
```

and converted to CBOR item:

```
[10, -3, 2, [
  41([
    [10, -4, 3],
    [-10, -11, 1]
  ]),
  3
]]
```

and finally to the binary string of:

```
0x840A220282D82982830A230383292A0103
```

A.7. Recursive Structure with Percent Encodings

This is a complex example having nested ARIs, some with percent-encoding needed. The human-friendly (but not valid URI) text for this case is:

```
ari:/adm10/rptt/rpt_with_param("text")
```

which is percent encoded to the real URI:

ari:/adm10/rppt/rpt_with_param(%22text%22)

which is translated to enumerated form:

ari:/10/-7/1(%22text%22)

and converted to CBOR item:

[10, -7, 1, ["text"]]

and finally to the binary string of:

0x840A2601816474657874

Authors' Addresses

Edward J. Birrane, III
The Johns Hopkins University Applied Physics Laboratory
11100 Johns Hopkins Rd.
Laurel, MD 20723
United States of America

Phone: [+1 443 778 7423](tel:+14437787423)
Email: Edward.Birrane@jhuapl.edu

Emery Annis
The Johns Hopkins University Applied Physics Laboratory

Email: Emery.Annis@jhuapl.edu

Brian Sipos
The Johns Hopkins University Applied Physics Laboratory

Email: brian.sipos+ietf@gmail.com