

Workgroup: CoRE Working Group
Internet-Draft: draft-ietf-core-href-12
Published: 6 March 2023
Intended Status: Standards Track
Expires: 7 September 2023
Authors: C. Bormann, Ed. H. Birkholz
 Universität Bremen TZI Fraunhofer SIT
 Constrained Resource Identifiers

Abstract

The Constrained Resource Identifier (CRI) is a complement to the Uniform Resource Identifier (URI) that serializes the URI components in Concise Binary Object Representation (CBOR) instead of a sequence of characters. This simplifies parsing, comparison and reference resolution in environments with severe limitations on processing power, code size, and memory size.

The present revision -12 of this draft adds a registry that is intended to provide full coverage for representing a URI scheme (and certain text strings used in their place) as negative integers.

About This Document

This note is to be removed before publishing as an RFC.

Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-ietf-core-href/>.

Discussion of this document takes place on the Constrained RESTful Environments Working Group mailing list (<mailto:core@ietf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/core/>. Subscribe at <https://www.ietf.org/mailman/listinfo/core/>.

Source for this draft and an issue tracker can be found at <https://github.com/core-wg/href>.

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 7 September 2023.

Copyright Notice

Copyright (c) 2023 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. Notational Conventions](#)
- [2. Constraints](#)
 - [2.1. Constraints not expressed by the data model](#)
- [3. Creation and Normalization](#)
- [4. Comparison](#)
- [5. CRI References](#)
 - [5.1. CBOR Serialization](#)
 - [5.1.1. The discard Section](#)
 - [5.1.2. Visualization](#)
 - [5.1.3. Examples](#)
 - [5.1.4. Specific Terminology](#)
 - [5.2. Ingesting and encoding a CRI Reference](#)
 - [5.3. Reference Resolution](#)
- [6. Relationship between CRIs, URIs and IRIs](#)
 - [6.1. Converting CRIs to URIs](#)
- [7. Extending CRIs](#)
 - [7.1. Extended CRI: Accommodating Percent Encoding \(PET\)](#)
- [8. Implementation Status](#)
- [9. Security Considerations](#)
- [10. IANA Considerations](#)
 - [10.1. CRI Scheme Numbers Registry](#)
 - [10.1.1. Instructions for the Designated Expert](#)
 - [10.1.2. Structure of Entries](#)
 - [10.1.3. Initial Registrations](#)
 - [10.2. Update to "Uniform Resource Identifier \(URI\) Schemes" Registry](#)

[11. References](#)

[11.1. Normative References](#)

[11.2. Informative References](#)

[Appendix A. Mapping Scheme Numbers to Scheme Names](#)

[Appendix B. The Small Print](#)

[Appendix C. Change Log](#)

[Acknowledgements](#)

[Contributors](#)

[Authors' Addresses](#)

1. Introduction

The [Uniform Resource Identifier \(URI\)](#) [RFC3986] and its most common usage, the URI reference, are the Internet standard for linking to resources in hypertext formats such as [HTML](#) [W3C.REC-html52-20171214] or the [HTTP "Link" header field](#) [RFC8288].

A URI reference is a sequence of characters chosen from the repertoire of US-ASCII characters. The individual components of a URI reference are delimited by a number of reserved characters, which necessitates the use of a character escape mechanism called "percent-encoding" when these reserved characters are used in a non-delimiting function. The resolution of URI references involves parsing a character sequence into its components, combining those components with the components of a base URI, merging path components, removing dot-segments, and recomposing the result back into a character sequence.

Overall, the proper handling of URI references is quite intricate. This can be a problem especially in [constrained environments](#) [RFC7228], where nodes often have severe code size and memory size limitations. As a result, many implementations in such environments support only an ad-hoc, informally-specified, bug-ridden, non-interoperable subset of half of RFC 3986.

This document defines the *Constrained Resource Identifier (CRI)* by constraining URIs to a simplified subset and serializing their components in [Concise Binary Object Representation \(CBOR\)](#) [RFC8949] instead of a sequence of characters. This allows typical operations on URI references such as parsing, comparison and reference resolution (including all corner cases) to be implemented in a comparatively small amount of code.

As a result of simplification, however, CRIs are not capable of expressing all URIs permitted by the generic syntax of RFC 3986 (hence the "constrained" in "Constrained Resource Identifier"). The supported subset includes all URIs of the [Constrained Application Protocol \(CoAP\)](#) [RFC7252], most URIs of the [Hypertext Transfer Protocol \(HTTP\)](#) [RFC9110], [Uniform Resource Names \(URNs\)](#) [RFC8141],

and other similar URIs. The exact constraints are defined in [Section 2](#).

1.1. Notational Conventions

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

In this specification, the term "byte" is used in its now customary sense as a synonym for "octet".

Terms defined in this document appear in *cursive* where they are introduced (rendered in plain text as the new term surrounded by underscores).

2. Constraints

A Constrained Resource Identifier consists of the same five components as a URI: scheme, authority, path, query, and fragment. The components are subject to the following constraints:

- C1.** The scheme name can be any Unicode string (see Definition D80 in [[Unicode](#)]) that matches the syntax of a URI scheme (see [Section 3.1](#) of [[RFC3986](#)], which constrains schemes to ASCII) and is lowercase (see Definition D139 in [[Unicode](#)]). The scheme is always present.
- C2.** An authority is always a host identified by an IP address or registered name, along with optional port information, and optionally preceded by user information.

Alternatively, the authority can be absent; the two cases for this defined in [Section 3.3](#) of [[RFC3986](#)] are modeled by two different values used in place of an absent authority:

- *the path can be root-based (zero or more path components that are each started in the URI with "/", as when the authority is present), or
- *the path can be rootless, which requires at least one path component.

(Note that in [Figure 1](#), no-authority is marked as a feature, as not all CRI implementations will support authority-less URIs.)

- C3.** A userinfo is a text string built out of unreserved characters ([Section 2.3](#) of [\[RFC3986\]](#)) or "sub-delims" ([Section 2.2](#) of [\[RFC3986\]](#)); any other character needs to be percent-encoded ([Section 7.1](#)). Note that this excludes the ":" character, which is commonly deprecated as a way to delimit a cleartext password in a userinfo.
- C4.** An IP address can be either an IPv4 address or an IPv6 address, optionally with a zone identifier [\[RFC6874\]](#). Future versions of IP are not supported (it is likely that a binary mapping would be strongly desirable, and that cannot be designed ahead of time, so these versions need to be added as a future extension if needed).
- C5.** A registered name is a sequence of one or more *labels*, which, when joined with dots (".") in between them, result in a Unicode string that is lowercase and in Unicode Normalization Form C (NFC) (see Definition D120 in [\[Unicode\]](#)). (The syntax may be further restricted by the scheme. As per [Section 3.2.2](#) of [\[RFC3986\]](#), a registered name can be empty, for which case a scheme can define a default for the host.)
- C6.** A port is always an integer in the range from 0 to 65535. Ports outside this range, empty ports (port subcomponents with no digits, see [Section 3.2.3](#) of [\[RFC3986\]](#)), or ports with redundant leading zeros, are not supported.
- C7.** The port is omitted if and only if the port would be the same as the scheme's default port (provided the scheme is defining such a default port) or the scheme is not using ports.
- C8.** A path consists of zero or more path segments. Note that a path of just a single zero-length path segment is allowed – this is considered equivalent to a path of zero path segments by HTTP and CoAP, but this equivalence does not hold for CRIs in general as they only perform normalization on the Syntax-Based Normalization level ([Section 6.2.2](#) of [\[RFC3986\]](#), not on the scheme-specific Scheme-Based Normalization level ([Section 6.2.3](#) of [\[RFC3986\]](#)).

(A CRI implementation may want to offer scheme-cognizant interfaces, performing this scheme-specific normalization for schemes it knows. The interface could assert which schemes the implementation knows and provide pre-normalized CRIs. This can also relieve the application from removing a lone zero-length

path segment before putting path segments into CoAP Options, i.e., from performing the check and jump in item 8 of [Section 6.4](#) of [\[RFC7252\]](#). See also [SP1](#) in [Appendix B.](#))

- C9.** A path segment can be any Unicode string that is in NFC, with the exception of the special "." and ".." complete path segments. Note that this includes the zero-length string.

If no authority is present in a CRI, the leading path segment cannot be empty. (See also [SP1](#) in [Appendix B.](#))

- C10.** A query always consists of one or more query parameters. A query parameter can be any Unicode string that is in NFC. It is often in the form of a "key=value" pair. When converting a CRI to a URI, query parameters are separated by an ampersand ("&") character. (This matches the structure and encoding of the target URI in CoAP requests.) Queries are optional; there is a difference between an absent query and a single query parameter that is the empty string.
- C11.** A fragment identifier can be any Unicode string that is in NFC. Fragment identifiers are optional; there is a difference between an absent fragment identifier and a fragment identifier that is the empty string.
- C12.** The syntax of registered names, path segments, query parameters, and fragment identifiers may be further restricted and sub-structured by the scheme. There is no support, however, for escaping sub-delimiters that are not intended to be used in a delimiting function.
- C13.** When converting a CRI to a URI, any character that is outside the allowed character range or is a delimiter in the URI syntax is percent-encoded. For CRIs, percent-encoding always uses the UTF-8 encoding form (see Definition D92 in [\[Unicode\]](#)) to convert the character to a sequence of bytes (that is then converted to a sequence of %HH triplets).

Examples for URIs at or beyond the boundaries of these constraints are in [SP2](#) in [Appendix B.](#)

2.1. Constraints not expressed by the data model

There are syntactically valid CRIs and CRI references that cannot be converted into a URI or URI reference, respectively.

For CRI references, this is acceptable -- they can be resolved still and result in a valid CRI that can be converted back. (An example of this is `[0, ["p"]]` which appends a slash and the path segment "p" to its base).

(Full) CRIs that do not correspond to a valid URI are not valid on their own, and cannot be used. Normatively they are characterized by the [Section 6.1](#) process producing a valid and syntax-normalized URI. For easier understanding, they are listed here:

- *CRIs (and CRI references) containing a path component "." or "..".

These would be removed by the `remove_dot_segments` algorithm of [\[RFC3986\]](#), and thus never produce a normalized URI after resolution.

(In CRI references, the discard value is used to afford segment removal, and with "." being an unreserved character, expressing them as "%2e" and "%2e%2e" is not even viable, let alone practical).

- *CRIs without authority whose path starts with two or more empty segments.

When converted to URIs, these would violate the requirement that in absence of an authority, a URI's path cannot begin with two slash characters, and they would be indistinguishable from a URI with a shorter path and a present but empty authority component.

- *CRIs without authority that are rootless and do not have a path component (e.g., ["a", true]), which would be indistinguishable from its root-based equivalent (["a"]) as both would have the URI a:.

3. Creation and Normalization

In general, resource identifiers are created on the initial creation of a resource with a certain resource identifier, or the initial exposition of a resource under a particular resource identifier.

A Constrained Resource Identifier **SHOULD** be created by the naming authority that governs the namespace of the resource identifier (see also [\[RFC8820\]](#)). For example, for the resources of an HTTP origin server, that server is responsible for creating the CRIs for those resources.

The naming authority **MUST** ensure that any CRI created satisfies the constraints defined in [Section 2](#). The creation of a CRI fails if the CRI cannot be validated to satisfy all of the constraints.

If a naming authority creates a CRI from user input, it **MAY** apply the following (and only the following) normalizations to get the CRI more likely to validate:

- *map the scheme name to lowercase ([C1](#));
- *map the registered name to NFC ([C5](#)) and split it on embedded dots;
- *elide the port if it is the default port for the scheme ([C7](#));
- *map path segments, query parameters and the fragment identifier to NFC form ([C9](#), [C10](#), [C11](#)).

Once a CRI has been created, it can be used and transferred without further normalization. All operations that operate on a CRI **SHOULD** rely on the assumption that the CRI is appropriately pre-normalized. (This does not contradict the requirement that when CRIs are transferred, recipients must operate on as-good-as untrusted input and fail gracefully in the face of malicious inputs.)

4. Comparison

One of the most common operations on CRIs is comparison: determining whether two CRIs are equivalent, without dereferencing the CRIs (using them to access their respective resource(s)).

Determination of equivalence or difference of CRIs is based on simple component-wise comparison. If two CRIs are identical component-by-component (using code-point-by-code-point comparison for components that are Unicode strings) then it is safe to conclude that they are equivalent.

This comparison mechanism is designed to minimize false negatives while strictly avoiding false positives. The constraints defined in [Section 2](#) imply the most common forms of syntax- and scheme-based normalizations in URIs, but do not comprise protocol-based normalizations that require accessing the resources or detailed knowledge of the scheme's dereference algorithm. False negatives can be caused, for example, by CRIs that are not appropriately pre-normalized and by resource aliases.

When CRIs are compared to select (or avoid) a network action, such as retrieval of a representation, fragment components (if any) should be excluded from the comparison.

5. CRI References

The most common usage of a Constrained Resource Identifier is to embed it in resource representations, e.g., to express a hyperlink

between the represented resource and the resource identified by the CRI.

This section defines the serialization of CRIs in [Concise Binary Object Representation \(CBOR\)](#) [RFC8949]. To reduce representation size, CRIs are not serialized directly. Instead, CRIs are indirectly referenced through *CRI references*. These take advantage of hierarchical locality and provide a very compact encoding. The CBOR serialization of CRI references is specified in [Section 5.1](#).

The only operation defined on a CRI reference is *reference resolution*: the act of transforming a CRI reference into a CRI. An application **MUST** implement this operation by applying the algorithm specified in [Section 5.3](#) (or any algorithm that is functionally equivalent to it).

The reverse operation of transforming a CRI into a CRI reference is unspecified; implementations are free to use any algorithm as long as reference resolution of the resulting CRI reference yields the original CRI. Notably, a CRI reference is not required to satisfy all of the constraints of a CRI; the only requirement on a CRI reference is that reference resolution **MUST** yield the original CRI.

When testing for equivalence or difference, applications **SHOULD NOT** directly compare CRI references; the references should be resolved to their respective CRI before comparison.

5.1. CBOR Serialization

A CRI or CRI reference is encoded as a CBOR array [RFC8949], with the structure as described in the [Concise Data Definition Language \(CDDL\)](#) [RFC8610] [including its control extensions](#) [RFC9165] as follows: RFC Ed.: throughout this section, please replace RFC-XXXX with the RFC number of this specification and remove this note.

```

; not expressed in this CDDL spec: trailing nulls to be left off

RFC-XXXX-Definitions = [CRI, CRI-Reference]

CRI = [
    scheme,
    authority / no-authority,
    local-part
]

CRI-Reference = [
    ((scheme / null, authority / no-authority)
     // discard),                ; relative reference
    local-part
]

local-part = (
    path / null,
    query / null,
    fragment / null
)

scheme      = scheme-name / scheme-id
scheme-name = text .regexp "[a-z][a-z0-9+.-]*)"
scheme-id   = nint

no-authority = NOAUTH-ROOTBASED / NOAUTH-ROOTLESS
NOAUTH-ROOTBASED = null .feature "no-authority"
NOAUTH-ROOTLESS = true .feature "no-authority"

authority    = [?userinfo, host, ?port]
userinfo     = (false, text .feature "userinfo")
host         = (host-ip // host-name)
host-name    = (*text) ; lowercase, NFC labels
host-ip      = (bytes .size 4 //
                (bytes .size 16, ?zone-id))
zone-id      = text
port         = 0..65535

discard      = DISCARD-ALL / 0..127
DISCARD-ALL  = true
path         = [*text]
query        = [*text]
fragment     = text

```

Figure 1: CDDL for CRI CBOR serialization

This CDDL specification is simplified for exposition and needs to be augmented by the following rule for interchange of CRIs and CRI references: Trailing null values **MUST** be removed, and two leading null values (scheme and authority both not given) are represented by using the discard alternative instead.

The rules scheme, authority, path, query, fragment correspond to the (sub-)components of a CRI, as described in [Section 2](#), with the addition of the discard section.

5.1.1. The discard Section


The discard section can be used in a CRI reference when neither a scheme nor an authority is present. It then expresses the operations performed on a base CRI by CRI references that are equivalent to URI references with relative paths and path prefixes such as `"/`, `"/.`, `"/..`, `"/../`, etc. `"/.` and `"/..` are not available in CRIs and are therefore expressed using discard after a normalization step, as is the presence or absence of a leading `"/`.

E.g., a simple URI reference `"foo"` specifies to remove one leading segment from the base URI's path, which is represented in the equivalent CRI reference discard section as the value 1; similarly `"/../foo"` removes two leading segments, represented as 2; and `"/foo"` removes all segments, represented in the discard section as the value true. The exact semantics of the section values are defined by [Section 5.3](#).

Most URI references that [Section 4.2](#) of [\[RFC3986\]](#) calls "relative references" (i.e., references that need to undergo a resolution process to obtain a URI) correspond to the CRI form that starts with discard. The exception are relative references with an authority (called a "network-path reference" in [Section 4.2](#) of [\[RFC3986\]](#)), which discard the entire path of the base CRI. These CRI references never carry a discard section: the value of discard defaults to true.

5.1.2. Visualization

The structure of a CRI reference is visualized using the somewhat limited means of a railroad diagram:



This visualization does not go into the details of the elements.

5.1.3. Examples

```
[-1,          / scheme -- equivalent to "coap" /  
[h'C6336401', / host /  
  61616],     / port /  
[".well-known", / path /  
  "core"]  
]
```

```
[true,          / discard /  
[".well-known", / path /  
  "core"],  
["rt=temperature-c"]] / query /
```

```
[-6,          / scheme -- equivalent to "did" /  
true,          / authority = NOAUTH-NOSLASH /  
["web:alice:bob"] / path /  
]
```

5.1.4. Specific Terminology

A CRI reference is considered *well-formed* if it matches the structure as expressed in [Figure 1](#) in CDDL, with the additional requirement that trailing null values are removed from the array.

A CRI reference is considered *absolute* if it is well-formed and the sequence of sections starts with a non-null scheme.

A CRI reference is considered *relative* if it is well-formed and the sequence of sections is empty or starts with a section other than those that would constitute a scheme.

5.2. Ingesting and encoding a CRI Reference

From an abstract point of view, a CRI Reference is a data structure with six sections:

scheme, authority, discard, path, query, fragment

Each of these sections can be unset ("null"), except for discard, which is always an unsigned number or true. If scheme and/or authority are non-null, discard must be true.

When ingesting a CRI Reference that is in the transfer form, those sections are filled in from the transfer form (unset sections are filled with null), and the following steps are performed:

*If the array is entirely empty, replace it with [0].

*If discard is present in the transfer form (i.e., the outer array starts with true or an unsigned number), set scheme and authority to null.

*If scheme and/or authority are present in the transfer form (i.e., the outer array starts with null, a text string, or a negative integer), set discard to true.

Upon encoding the abstract form into the transfer form, the inverse processing is performed: If scheme and/or authority are not null, the discard value is not transferred (it must be true in this case). If they are both null, they are both left out and only discard is transferred. Trailing null values are removed from the array. As a special case, an empty array is sent in place for a remaining [0] (URI "").

5.3. Reference Resolution

The term "relative" implies that a "base CRI" exists against which the relative reference is applied. Aside from fragment-only references, relative references are only usable when a base CRI is known.

The following steps define the process of resolving any well-formed CRI reference against a base CRI so that the result is a CRI in the form of an absolute CRI reference:

1. Establish the base CRI of the CRI reference and express it in the form of an abstract absolute CRI reference.
2. Initialize a buffer with the sections from the base CRI.
3. If the value of discard is true in the CRI reference (which is implicitly the case when scheme and/or authority are present in the reference), replace the path in the buffer with the empty array, unset query and fragment, and set a true authority to null. If the value of discard is an unsigned number, remove as many elements from the end of the path array; if it is non-zero, unset query and fragment.

Set discard to true in the buffer.

4. If the path section is set in the CRI reference, append all elements from the path array to the array in the path section in the buffer; unset query and fragment.
5. Apart from the path and discard, copy all non-null sections from the CRI reference to the buffer in sequence; unset fragment in the buffer if query is non-null in the CRI reference (and therefore has been copied to the buffer).

6. Return the sections in the buffer as the resolved CRI.

6. Relationship between CRIs, URIs and IRIs

CRIs are meant to replace both [Uniform Resource Identifiers \(URIs\)](#) [RFC3986] and [Internationalized Resource Identifiers \(IRIs\)](#) [RFC3987] in [constrained environments](#) [RFC7228]. Applications in these environments may never need to use URIs and IRIs directly, especially when the resource identifier is used simply for identification purposes or when the CRI can be directly converted into a CoAP request.

However, it may be necessary in other environments to determine the associated URI or IRI of a CRI, and vice versa. Applications can perform these conversions as follows:

CRI to URI

A CRI is converted to a URI as specified in [Section 6.1](#).

URI to CRI

The method of converting a URI to a CRI is unspecified; implementations are free to use any algorithm as long as converting the resulting CRI back to a URI yields an equivalent URI.

CRI to IRI

A CRI can be converted to an IRI by first converting it to a URI as specified in [Section 6.1](#), and then converting the URI to an IRI as described in [Section 3.2](#) of [RFC3987].

IRI to CRI

An IRI can be converted to a CRI by first converting it to a URI as described in [Section 3.1](#) of [RFC3987], and then converting the URI to a CRI as described above.

Everything in this section also applies to CRI references, URI references and IRI references.

6.1. Converting CRIs to URIs

Applications **MUST** convert a CRI reference to a URI reference by determining the components of the URI reference according to the following steps and then recomposing the components to a URI reference string as specified in [Section 5.3](#) of [RFC3986].

scheme

If the CRI reference contains a scheme section, the scheme component of the URI reference consists of the value of that section, if text (scheme-name); or, if a negative integer is given (scheme-id), the lower case scheme name corresponding to

the scheme number as per the CRI Scheme Numbers registry [Section 10.1](#). Otherwise, the scheme component is unset.

authority

If the CRI reference contains a host-name or host-ip item, the authority component of the URI reference consists of a host subcomponent, optionally followed by a colon (":") character and a port subcomponent, optionally preceded by a userinfo subcomponent. Otherwise, the authority component is unset.

The host subcomponent consists of the value of the host-name or host-ip item.

The userinfo subcomponent, if present, is turned into a single string by appending a "@". Otherwise, both the subcomponent and the "@" sign are omitted. Any character in the value of the userinfo elements that is not in the set of unreserved characters ([Section 2.3](#) of [[RFC3986](#)]) or "sub-delims" ([Section 2.2](#) of [[RFC3986](#)]) **MUST** be percent-encoded.

The host-name is turned into a single string by joining the elements separated by dots ("."). Any character in the elements of a host-name item that is a dot ("."), or not in the set of unreserved characters ([Section 2.3](#) of [[RFC3986](#)]) or "sub-delims" ([Section 2.2](#) of [[RFC3986](#)]) **MUST** be percent-encoded.

The value of a host-ip item **MUST** be represented as a string that matches the "IPv4address" or "IP-literal" rule ([Section 3.2.2](#) of [[RFC3986](#)]). Any zone-id is appended to the string, separated by "%25" as defined in [Section 2](#) of [[RFC6874](#)], or as specified in a superseding zone-id specification document [[I-D.carpenter-6man-rfc6874bis](#)]; this also leads to a modified "IP-literal" rule as specified in these documents.

If the CRI reference contains a port item, the port subcomponent consists of the value of that item in decimal notation. Otherwise, the colon (":") character and the port subcomponent are both omitted.

path

If the CRI reference contains a discard item of value true, the path component is considered *rooted*. If it contains a discard item of value 0 and the path item is present, the conversion fails. If it contains a positive discard item, the path component is considered *unrooted* and prefixed by as many "../" components as the discard value minus one indicates. If the discard value is 1 and the first element of the path contains a :, the path component is prefixed by "./" (this avoids the first element to

appear as supplying a URI scheme; compare path-noscheme in [Section 4.2](#) of [\[RFC3986\]](#)).

If the discard item is not present and the CRI reference contains an authority that is true, the path component of the URI reference is considered unrooted. Otherwise, the path component is considered rooted.

If the CRI reference contains one or more path items, the path component is constructed by concatenating the sequence of representations of these items. These representations generally contain a leading slash ("/") character and the value of each item, processed as discussed below. The leading slash character is omitted for the first path item only if the path component is considered "unrooted".

Any character in the value of a path item that is not in the set of unreserved characters or "sub-delims" or a colon (":") or commercial at ("@") character **MUST** be percent-encoded.

If the authority component is present (not null or true) and the path component does not match the "path-abempty" rule ([Section 3.3](#) of [\[RFC3986\]](#)), the conversion fails.

If the authority component is not present, but the scheme component is, and the path component does not match the "path-absolute", "path-rootless" (authority == true) or "path-empty" rule ([Section 3.3](#) of [\[RFC3986\]](#)), the conversion fails.

If neither the authority component nor the scheme component are present, and the path component does not match the "path-absolute", "path-noscheme" or "path-empty" rule ([Section 3.3](#) of [\[RFC3986\]](#)), the conversion fails.

query

If the CRI reference contains one or more query items, the query component of the URI reference consists of the value of each item, separated by an ampersand ("&") character. Otherwise, the query component is unset.

Any character in the value of a query item that is not in the set of unreserved characters or "sub-delims" or a colon (":"), commercial at ("@"), slash ("/") or question mark ("?") character **MUST** be percent-encoded. Additionally, any ampersand character ("&") in the item value **MUST** be percent-encoded.

fragment

If the CRI reference contains a fragment item, the fragment component of the URI reference consists of the value of that item. Otherwise, the fragment component is unset.

Any character in the value of a fragment item that is not in the set of unreserved characters or "sub-delims" or a colon (":"), commercial at ("@"), slash ("/") or question mark ("?") character **MUST** be percent-encoded.

7. Extending CRIs

CRIs have been designed to relieve implementations operating on CRIs from string scanning, which both helps constrained implementations and implementations that need to achieve high throughput.

The CRI structure described up to this point is termed the *Basic CRI*. It should be sufficient for all applications that use the CoAP protocol, as well as most other protocols employing URIs.

However, Basic CRIs have one limitation: They do not support URI components that *require* percent-encoding ([Section 2.1](#) of [\[RFC3986\]](#)) to represent them in the URI syntax, except where that percent-encoding is used to escape the main delimiter in use.

E.g., the URI

`https://alice/3%2f4-inch`

is represented by the basic CRI

`[-4, ["alice"], ["3/4-inch"]]`

However, percent-encoding that is used at the application level is not supported by basic CRIs:

`did:web:alice:7%3A1-balun`

Extended forms of CRIs may be defined to enable these applications. They will generally extend the potential values of text components of URIs, such as userinfo, hostnames, paths, queries, and fragments.

One such extended form is described in the following [Section 7.1](#). Consumers of CRIs will generally notice when an extended form is in use, by finding structures that do not match the CDDL rules given in [Figure 1](#). Future definitions of extended forms need to strive to be distinguishable in their structures from the extended form presented here as well as other future forms.

7.1. Extended CRI: Accommodating Percent Encoding (PET)

This section presents a method to represent percent-encoded segments of userinfo, hostnames, paths, and queries, as well as fragments.

The four CDDL rules

```

userinfo    = (false, text .feature "userinfo")
host-name   = (*text)
path        = [*text]
query       = [*text]
fragment    = text

```

are replaced with

```

userinfo    = (false, text-or-pet .feature "userinfo")
host-name   = (*text-or-pet)
path        = [*text-or-pet]
query       = [*text-or-pet]
fragment    = text-or-pet

```

```

text-or-pet = text /
    text-pet-sequence .feature "extended-cri"

```

```

; text1 and pet1 alternating, at least one pet1:
text-pet-sequence = [?text1, ((+(pet1, text1), ?pet1) // pet1)]
; pet is percent-encoded bytes
pet1 = bytes .ne ''
text1 = text .ne ""

```

That is, for each of the host-name, path, and query segments, and for the userinfo and fragment components, an alternate representation is provided besides a simple text string: a non-empty array of alternating non-blank text and byte strings, the text strings of which stand for non-percent-encoded text, while the byte strings retain the special semantics of percent-encoded text without actually being percent-encoded.

The above DID URI can now be represented as:

```
[-6, true, [{"web:alice:7", ':', "1-balun"}]]
```

Unlike the text elements of a path or a query, which through CoAP's heritage are designed to be processable element by element, a text-pet-sequence does not usually produce semantically meaningful array elements. This consequence of the flexibility in delimiters offered in URIs is demonstrated by this example, which structurally singles out the one ':' that is *not* a delimiter at the application level. Applications designed for using CRIs will generally avoid using the extended-cri feature. Applications using existing URI structures that require text-pet-sequence elements for their representation typically need to process them byte by byte.

8. Implementation Status

With the exception of the authority=true fix, host-names split into labels, and [Section 7.1](https://gitlab.com/), CRIs are implemented in <https://gitlab.com/>

chrysn/micrurus. A goolang implementation of version -10 of this document is found at: <https://github.com/thomas-fossati/href>

9. Security Considerations

Parsers of CRI references must operate on input that is assumed to be untrusted. This means that parsers **MUST** fail gracefully in the face of malicious inputs. Additionally, parsers **MUST** be prepared to deal with resource exhaustion (e.g., resulting from the allocation of big data items) or exhaustion of the call stack (stack overflow). See [Section 10](#) of [[RFC8949](#)] for additional security considerations relating to CBOR.

The security considerations discussed in [Section 7](#) of [[RFC3986](#)] and [Section 8](#) of [[RFC3987](#)] for URIs and IRIs also apply to CRIs.

10. IANA Considerations

10.1. CRI Scheme Numbers Registry

This specification defines a new "CRI Scheme Numbers" sub-registry in the "CoRE Parameters" registry [[IANA.core-parameters](#)], with the policy "Expert Review" ([Section 4.5](#) of [[BCP26](#)]). The objective is to have CRI scheme number values registered for all registered URI schemes (Uniform Resource Identifier (URI) Schemes registry), as well as exceptionally for certain text strings that the Designated Expert considers widely used in constrained applications in place of URI scheme names.

10.1.1. Instructions for the Designated Expert

The expert is instructed to be frugal in the allocation of CRI values with short representations (1+0 and 1+1 encoding), keeping them in reserve for applications that are likely to enjoy wide use and can make good use of their shortness.

When the expert notices that a registration has been made in the Uniform Resource Identifier (URI) Schemes registry (see also [Section 10.2](#)), the expert is requested to initiate a parallel registration in the CRI Scheme Numbers registry. CRI values in the range between 1000 and 20000 (inclusive) should be assigned unless a shorter representation in CRIs appears desirable.

The expert exceptionally also may make such a registration for text strings that have not been registered in the Uniform Resource Identifier (URI) Schemes registry if and only if the expert considers the to be in wide use in place of URI scheme names in constrained applications. (Note that the initial registrations in [Table 1](#) already include such registrations for the text strings "mqtt" and "mqttp".)

A registration in the CRI Scheme Numbers registry does not imply that a URI scheme under this name exists or has been registered in the Uniform Resource Identifier (URI) Schemes registry -- it essentially is only providing an integer identifier for an otherwise uninterpreted text string.

Any questions or issues that might interest a wider audience might be raised by the expert on the core-parameters@ietf.org mailing list for a time-limited discussion.

10.1.2. Structure of Entries

Each entry in the registry must include:

CRI value:

A negative integer unique in this registry

URI scheme name:

a text string that would be acceptable for registration as a URI Scheme Name in the Uniform Resource Identifier (URI) Schemes registry

Reference:

a reference to a document, if available, or the registrant

10.1.3. Initial Registrations

The initial registrations for the CRI Scheme Numbers registry are provided in [Table 1](#).

10.2. Update to "Uniform Resource Identifier (URI) Schemes" Registry

[[RFC7595](#)] is updated to add the following note in the "Uniform Resource Identifier (URI) Schemes" Registry [[IANA.uri-schemes](#)]:

The CRI Scheme Numbers Registry registers numeric identifiers for what essentially are URI Scheme names. Registrants for the Uniform Resource Identifier (URI) Schemes Registry are requested to make a parallel registration in the CRI Scheme Numbers registry. The number for this registration will be assigned by the Designated Expert for that registry.

11. References

11.1. Normative References

- [[BCP26](#)] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/rfc/rfc8126>>.

[I-D.carpenter-6man-rfc6874bis]

Carpenter, B. E., Cheshire, S., and R. M. Hinden, "Representing IPv6 Zone Identifiers in Address Literals and Uniform Resource Identifiers", Work in Progress, Internet-Draft, draft-carpenter-6man-rfc6874bis-03, 8 February 2022, <<https://datatracker.ietf.org/doc/html/draft-carpenter-6man-rfc6874bis-03>>.

[IANA.core-parameters] IANA, "Constrained RESTful Environments (CoRE) Parameters", <<https://www.iana.org/assignments/core-parameters>>.

[IANA.uri-schemes] IANA, "Uniform Resource Identifier (URI) Schemes", <<https://www.iana.org/assignments/uri-schemes>>.

[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/rfc/rfc2119>>.

[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/rfc/rfc3986>>.

[RFC3987] Duerst, M. and M. Suignard, "Internationalized Resource Identifiers (IRIs)", RFC 3987, DOI 10.17487/RFC3987, January 2005, <<https://www.rfc-editor.org/rfc/rfc3987>>.

[RFC6874] Carpenter, B., Cheshire, S., and R. Hinden, "Representing IPv6 Zone Identifiers in Address Literals and Uniform Resource Identifiers", RFC 6874, DOI 10.17487/RFC6874, February 2013, <<https://www.rfc-editor.org/rfc/rfc6874>>.

[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/rfc/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/rfc/rfc8174>>.

[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/rfc/rfc8610>>.

[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/rfc/rfc8949>>.

[RFC9165]

Bormann, C., "Additional Control Operators for the Concise Data Definition Language (CDDL)", RFC 9165, DOI 10.17487/RFC9165, December 2021, <<https://www.rfc-editor.org/rfc/rfc9165>>.

[Unicode]

The Unicode Consortium, "The Unicode Standard, Version 13.0.0", ISBN 978-1-936213-26-9, March 2020, <<https://www.unicode.org/versions/Unicode13.0.0/>>.

11.2. Informative References

[RFC4180]

Shafranovich, Y., "Common Format and MIME Type for Comma-Separated Values (CSV) Files", RFC 4180, DOI 10.17487/RFC4180, October 2005, <<https://www.rfc-editor.org/rfc/rfc4180>>.

[RFC7228]

Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<https://www.rfc-editor.org/rfc/rfc7228>>.

[RFC7252]

Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/rfc/rfc7252>>.

[RFC8141]

Saint-Andre, P. and J. Klensin, "Uniform Resource Names (URNs)", RFC 8141, DOI 10.17487/RFC8141, April 2017, <<https://www.rfc-editor.org/rfc/rfc8141>>.

[RFC8288]

Nottingham, M., "Web Linking", RFC 8288, DOI 10.17487/RFC8288, October 2017, <<https://www.rfc-editor.org/rfc/rfc8288>>.

[RFC8820]

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

[RFC9110]

Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP Semantics", STD 97, RFC 9110, DOI 10.17487/

RFC9110, June 2022, <<https://www.rfc-editor.org/rfc/rfc9110>>.

[W3C.REC-htm152-20171214] Danilo, A., Ed., Eicholz, A., Ed., Moon, S., Ed., Faulkner, S., Ed., and T. Leithead, Ed., "HTML 5.2", W3C REC REC-htm152-20171214, W3C REC-htm152-20171214, 14 December 2017, <<https://www.w3.org/TR/2017/REC-htm152-20171214/>>.

Appendix A. Mapping Scheme Numbers to Scheme Names

[Table 1](#) defines the initial mapping from CRI scheme numbers to URI scheme names.

CRI value	URI scheme	Reference
-1	coap	[RFCthis]
-2	coaps	[RFCthis]
-3	http	[RFCthis]
-4	https	[RFCthis]
-5	urn	[RFCthis]
-6	did	[RFCthis]
-7	coap+tcp	[RFCthis]
-8	coaps+tcp	[RFCthis]
-9	coap+ws	[RFCthis]
-10	coaps+ws	[RFCthis]
-1025	telnet	[RFCthis]
-1046	ldap	[RFCthis]
-1056	ms-virtualtouchpad	[RFCthis]
-1091	fax	[RFCthis]
-1107	ves	[RFCthis]
-1147	submit	[RFCthis]
-1192	gg	[RFCthis]
-1219	simplex	[RFCthis]
-1240	ms-settings-nfctransactions	[RFCthis]
-1241	secret-token	[RFCthis]
-1249	acap	[RFCthis]
-1276	openpgp4fpr	[RFCthis]
-1300	ms-mixedrealitycapture	[RFCthis]
-1307	ymsg	[RFCthis]
-1320	iris.xpcs	[RFCthis]
-1351	turns	[RFCthis]
-1367	opaquelocktoken	[RFCthis]
-1499	platform	[RFCthis]
-1597	sftp	[RFCthis]
-1613	vscode	[RFCthis]
-1649	mqtt	[RFCthis]

CRI value	URI scheme	Reference
-1664	ms-settings	[RFCthis]
-1690	doi	[RFCthis]
-1720	file	[RFCthis]
-1729	dvb	[RFCthis]
-1760	magnet	[RFCthis]
-1768	calculator	[RFCthis]
-1836	ssh	[RFCthis]
-1966	gopher	[RFCthis]
-1985	ms-gamingoverlay	[RFCthis]
-1997	z39.50	[RFCthis]
-2032	ms-secondary-screen-setup	[RFCthis]
-2038	fido	[RFCthis]
-2085	mumble	[RFCthis]
-2095	ms-settings-cloudstorage	[RFCthis]
-2106	imap	[RFCthis]
-2152	ms-officeapp	[RFCthis]
-2233	pwid	[RFCthis]
-2236	drm	[RFCthis]
-2264	tag	[RFCthis]
-2369	feed	[RFCthis]
-2460	ipps	[RFCthis]
-2484	xmlrpc.beeps	[RFCthis]
-2492	jms	[RFCthis]
-2542	wpid	[RFCthis]
-2669	barion	[RFCthis]
-2675	onenote	[RFCthis]
-2695	icon	[RFCthis]
-2769	message	[RFCthis]
-2800	ms-enrollment	[RFCthis]
-2804	bolo	[RFCthis]
-2817	diaspora	[RFCthis]
-2833	microsoft.windows.camera.picker	[RFCthis]
-2864	notes	[RFCthis]
-2866	amss	[RFCthis]
-2873	tip	[RFCthis]
-3018	fm	[RFCthis]
-3042	rtmfp	[RFCthis]
-3060	reload	[RFCthis]
-3111	pres	[RFCthis]
-3232	acd	[RFCthis]
-3362	prospero	[RFCthis]
-3364	geo	[RFCthis]
-3414	snmp	[RFCthis]
-3483	iris.beep	[RFCthis]
-3510	maps	[RFCthis]

CRI value	URI scheme	Reference
-3575	content	[RFCthis]
-3618	pack	[RFCthis]
-3619	keyparc	[RFCthis]
-3632	mongodb	[RFCthis]
-3693	smb	[RFCthis]
-3796	graph	[RFCthis]
-3818	filesystem	[RFCthis]
-3839	payment	[RFCthis]
-3840	ms-settings-bluetooth	[RFCthis]
-3951	palm	[RFCthis]
-4027	hyper	[RFCthis]
-4043	microsoft.windows.camera	[RFCthis]
-4067	mvn	[RFCthis]
-4098	mtqp	[RFCthis]
-4130	jabber	[RFCthis]
-4275	mms	[RFCthis]
-4343	skype	[RFCthis]
-4351	oid	[RFCthis]
-4420	dict	[RFCthis]
-4454	attachment	[RFCthis]
-4662	ocf	[RFCthis]
-4807	isostore	[RFCthis]
-4816	redis	[RFCthis]
-4862	ms-settings-privacy	[RFCthis]
-4877	ms-settings-wifi	[RFCthis]
-5004	v-event	[RFCthis]
-5020	com-eventbrite-attendee	[RFCthis]
-5105	teliaeid	[RFCthis]
-5222	itms	[RFCthis]
-5234	fish	[RFCthis]
-5285	dtn	[RFCthis]
-5298	vscode-insiders	[RFCthis]
-5304	tftp	[RFCthis]
-5347	rtsp	[RFCthis]
-5358	adiumextra	[RFCthis]
-5464	smp	[RFCthis]
-5470	ms-eyecontrolspeech	[RFCthis]
-5479	ms-settings-language	[RFCthis]
-5491	mqtt	[RFCthis]
-5595	wyciwyg	[RFCthis]
-5596	hcp	[RFCthis]
-5619	go	[RFCthis]
-5673	rediss	[RFCthis]
-5683	ms-settings-cellular	[RFCthis]
-5743	ldaps	[RFCthis]

CRI value	URI scheme	Reference
-5843	z39.50s	[RFCthis]
-5886	bitcoincash	[RFCthis]
-5960	ms-mobileplans	[RFCthis]
-6182	pttp	[RFCthis]
-6208	facetime	[RFCthis]
-6289	gtalk	[RFCthis]
-6348	afp	[RFCthis]
-6361	mss	[RFCthis]
-6426	ms-settings-notifications	[RFCthis]
-6448	psyc	[RFCthis]
-6488	tv	[RFCthis]
-6514	wifi	[RFCthis]
-6523	sarif	[RFCthis]
-6539	moz	[RFCthis]
-6659	ms-lockscreencomponent-config	[RFCthis]
-6716	cabal	[RFCthis]
-6734	ms-media-stream-id	[RFCthis]
-6780	mupdate	[RFCthis]
-6793	dis	[RFCthis]
-6804	nih	[RFCthis]
-6809	ms-help	[RFCthis]
-6909	soap.beep	[RFCthis]
-6998	iotdisco	[RFCthis]
-7027	acr	[RFCthis]
-7040	ms-newsandinterests	[RFCthis]
-7089	hxxp	[RFCthis]
-7096	ms-settings-location	[RFCthis]
-7125	soap.beeps	[RFCthis]
-7301	ipn	[RFCthis]
-7309	nntp	[RFCthis]
-7316	query	[RFCthis]
-7334	smtp	[RFCthis]
-7335	ms-spd	[RFCthis]
-7400	ni	[RFCthis]
-7403	ms-excel	[RFCthis]
-7421	ms-settings-power	[RFCthis]
-7435	pop	[RFCthis]
-7447	session	[RFCthis]
-7582	ms-infopath	[RFCthis]
-7701	ms-word	[RFCthis]
-7715	web+ap	[RFCthis]
-7791	steam	[RFCthis]
-7995	cstr	[RFCthis]
-8008	web3	[RFCthis]
-8064	videotex	[RFCthis]

CRI value	URI scheme	Reference
-8069	nfs	[RFCthis]
-8094	udp	[RFCthis]
-8102	ed2k	[RFCthis]
-8138	ms-getoffice	[RFCthis]
-8203	sgn	[RFCthis]
-8331	data	[RFCthis]
-8364	swidpath	[RFCthis]
-8385	fuchsia-pkg	[RFCthis]
-8395	ms-screensketch	[RFCthis]
-8426	hxxps	[RFCthis]
-8487	unreal	[RFCthis]
-8555	ens	[RFCthis]
-8585	ms-settings-camera	[RFCthis]
-8619	stun	[RFCthis]
-8673	ms-stickers	[RFCthis]
-8775	spotify	[RFCthis]
-8860	starknet	[RFCthis]
-8890	ms-settings-emailandaccounts	[RFCthis]
-8907	market	[RFCthis]
-8967	ms-powerpoint	[RFCthis]
-9001	rtsp	[RFCthis]
-9064	p1	[RFCthis]
-9128	aw	[RFCthis]
-9132	mailserver	[RFCthis]
-9186	irc6	[RFCthis]
-9338	ms-settings-lock	[RFCthis]
-9339	hcap	[RFCthis]
-9350	drop	[RFCthis]
-9419	icap	[RFCthis]
-9437	xcon-userid	[RFCthis]
-9457	leapfrogans	[RFCthis]
-9461	ipfs	[RFCthis]
-9479	bitcoin	[RFCthis]
-9555	apt	[RFCthis]
-9605	ms-whiteboard-cmd	[RFCthis]
-9669	ssb	[RFCthis]
-9725	aaas	[RFCthis]
-9734	ar	[RFCthis]
-9767	proxy	[RFCthis]
-9773	res	[RFCthis]
-9780	msrps	[RFCthis]
-9795	aim	[RFCthis]
-9826	tool	[RFCthis]
-9842	finger	[RFCthis]
-9900	turn	[RFCthis]

CRI value	URI scheme	Reference
-9901	num	[RFCthis]
-9903	svn	[RFCthis]
-9904	ut2004	[RFCthis]
-9932	ms-visio	[RFCthis]
-10008	eid	[RFCthis]
-10100	wss	[RFCthis]
-10103	gizmo-project	[RFCthis]
-10172	dlna-playsingle	[RFCthis]
-10224	swh	[RFCthis]
-10337	dat	[RFCthis]
-10348	cap	[RFCthis]
-10355	z39.50r	[RFCthis]
-10412	xcon	[RFCthis]
-10430	gitoid	[RFCthis]
-10524	hydrazone	[RFCthis]
-10565	example	[RFCthis]
-10699	crid	[RFCthis]
-10717	teamspeak	[RFCthis]
-10743	elsi	[RFCthis]
-10769	dtmi	[RFCthis]
-10840	ftp	[RFCthis]
-10902	ms-drive-to	[RFCthis]
-10903	upt	[RFCthis]
-10911	appdata	[RFCthis]
-11039	callto	[RFCthis]
-11131	ms-remotedesktop-launch	[RFCthis]
-11139	dweb	[RFCthis]
-11264	lastfm	[RFCthis]
-11307	xmlrpc.beep	[RFCthis]
-11342	ms-whiteboard	[RFCthis]
-11465	first-run-pen-experience	[RFCthis]
-11473	webcal	[RFCthis]
-11553	adt	[RFCthis]
-11566	vemmi	[RFCthis]
-11590	cvs	[RFCthis]
-11629	taler	[RFCthis]
-11688	ms-inputapp	[RFCthis]
-11864	git	[RFCthis]
-11893	irc	[RFCthis]
-11936	ms-settings-workplace	[RFCthis]
-12171	blob	[RFCthis]
-12173	modem	[RFCthis]
-12188	msnim	[RFCthis]
-12268	iris.lwz	[RFCthis]
-12302	ms-sttoverlay	[RFCthis]

CRI value	URI scheme	Reference
-12321	lbry	[RFCthis]
-12334	rmi	[RFCthis]
-12346	ms-restoretacompanion	[RFCthis]
-12482	ms-useractivityset	[RFCthis]
-12485	dab	[RFCthis]
-12491	about	[RFCthis]
-12500	embedded	[RFCthis]
-12501	rtmp	[RFCthis]
-12526	ircs	[RFCthis]
-12558	mid	[RFCthis]
-12573	sip	[RFCthis]
-12593	ipns	[RFCthis]
-12666	dvx	[RFCthis]
-12706	android	[RFCthis]
-12747	wtai	[RFCthis]
-12831	ms-search-repair	[RFCthis]
-12838	microsoft.windows.camera.multipicker	[RFCthis]
-12857	ms-settings-screenrotation	[RFCthis]
-12879	rtspu	[RFCthis]
-12914	ms-screenclip	[RFCthis]
-12943	aaa	[RFCthis]
-12954	xmpp	[RFCthis]
-12988	soldat	[RFCthis]
-13041	lorawan	[RFCthis]
-13054	beshare	[RFCthis]
-13077	sips	[RFCthis]
-13081	iris.xpc	[RFCthis]
-13113	simpleledger	[RFCthis]
-13127	vsis	[RFCthis]
-13207	matrix	[RFCthis]
-13307	otpauth	[RFCthis]
-13336	cid	[RFCthis]
-13352	service	[RFCthis]
-13417	h323	[RFCthis]
-13438	ms-settings-connectabledevices	[RFCthis]
-13452	payto	[RFCthis]
-13463	ms-settings-displays-topology	[RFCthis]
-13505	lvlt	[RFCthis]
-13596	ms-walk-to	[RFCthis]
-13672	dns	[RFCthis]
-13730	quic-transport	[RFCthis]
-13762	paparazzi	[RFCthis]
-13766	ms-people	[RFCthis]
-13889	xri	[RFCthis]
-13894	onenote-cmd	[RFCthis]

CRI value	URI scheme	Reference
-13934	dav	[RFCthis]
-14003	content-type	[RFCthis]
-14068	sms	[RFCthis]
-14119	ms-publisher	[RFCthis]
-14197	xfire	[RFCthis]
-14250	secondlife	[RFCthis]
-14260	ark	[RFCthis]
-14301	iax	[RFCthis]
-14312	msrp	[RFCthis]
-14475	swid	[RFCthis]
-14590	tn3270	[RFCthis]
-14596	ms-appinstaller	[RFCthis]
-14627	stuns	[RFCthis]
-14688	dpp	[RFCthis]
-14701	ms-secondary-screen-controller	[RFCthis]
-14764	browserext	[RFCthis]
-14820	chrome	[RFCthis]
-14878	pkcs11	[RFCthis]
-15066	dlna-playcontainer	[RFCthis]
-15155	spiffe	[RFCthis]
-15207	uuid-in-package	[RFCthis]
-15261	ms-settings-proximity	[RFCthis]
-15356	things	[RFCthis]
-15377	ms-gamebarservices	[RFCthis]
-15379	shc	[RFCthis]
-15547	ipp	[RFCthis]
-15552	mailto	[RFCthis]
-15558	ms-browser-extension	[RFCthis]
-15838	shttp (OBSOLETE)	[RFCthis]
-15842	acct	[RFCthis]
-15849	w3	[RFCthis]
-15869	wais	[RFCthis]
-15928	qb	[RFCthis]
-15947	ms-search	[RFCthis]
-16043	ms-settings-airplanemode	[RFCthis]
-16045	jar	[RFCthis]
-16069	tel	[RFCthis]
-16074	dntp	[RFCthis]
-16160	chrome-extension	[RFCthis]
-16193	cast	[RFCthis]
-16326	view-source	[RFCthis]
-16356	im	[RFCthis]
-16358	resource	[RFCthis]
-16378	ms-calculator	[RFCthis]
-16380	news	[RFCthis]

CRI value	URI scheme	Reference
-16415	wcr	[RFCthis]
-16523	casts	[RFCthis]
-16689	ms-access	[RFCthis]
-16723	grd	[RFCthis]
-16750	rsync	[RFCthis]
-16773	lpa	[RFCthis]
-16850	afs	[RFCthis]
-16874	bb	[RFCthis]
-16884	ham	[RFCthis]
-16926	info	[RFCthis]
-16972	ms-meetnow	[RFCthis]
-17117	ms-project	[RFCthis]
-17172	ethereum	[RFCthis]
-17225	thismessage	[RFCthis]
-17226	vnc	[RFCthis]
-17232	snews	[RFCthis]
-17245	sieve	[RFCthis]
-17269	feedready	[RFCthis]
-17271	mt	[RFCthis]
-17288	ws	[RFCthis]
-17338	ms-transit-to	[RFCthis]
-17346	ventrilo	[RFCthis]
-17357	iris	[RFCthis]

Table 1: Mapping Scheme Numbers to Scheme Names

The assignments from this table can be extracted from the XML form of this document (when stored in a file "this.xml") into CSV form [RFC4180] using this short Ruby program:

```
require 'rexml/document'; include REXML
XPath.each(Document.new(File.read("this.xml")), "///tr") {|row|
  puts XPath.each(row, "td").map{|d|d.text()}[0..1].join(",")}
```

Appendix B. The Small Print

This appendix lists a few corner cases of URI semantics that implementers of CRIs need to be aware of, but that are not representative of the normal operation of CRIs.

SP1. Initial (Lone/Leading) Empty Path Segments:

**Lone empty path segments:* As per [RFC3986], s://x is distinct from s://x/ -- i.e., a URI with an empty path is different from one with a lone empty path segment. However, in HTTP, CoAP, they are implicitly aliased (for CoAP, in item 8 of [Section 6.4](#) of [RFC7252]). As per item 7 of [Section 6.5](#) of [RFC7252], recomposition of a URI without Uri-Path Options from the other

URI-related CoAP Options produces `s://x/`, not `s://x --` CoAP prefers the lone empty path segment form. TBD: add similar text for HTTP, if that can be made. After discussing HTTP semantics, [Section 6.2.3](#) of [\[RFC3986\]](#) even states:

In general, a URI that uses the generic syntax for authority with an empty path should be normalized to a path of `"/`.

**Leading empty path segments without authority:* Somewhat related, note also that URIs and URI references that do not carry an authority cannot represent initial empty path segments (i.e., that are followed by further path segments): `s://x//foo` works, but in a `s://foo` URI or an (absolute-path) URI reference of the form `//foo` the double slash would be mis-parsed as leading in to an authority.

SP2. Constraints ([Section 2](#)) of CRIs/basic CRIs

While most URIs in everyday use can be converted to CRIs and back to URIs matching the input after syntax-based normalization of the URI, these URIs illustrate the constraints by example:

**https://host%ffname, https://example.com/x?data=%ff*

All URI components must, after percent decoding, be valid UTF-8 encoded text. Bytes that are not valid UTF-8 show up, for example, in BitTorrent web seeds.

**https://example.com/component%3bone;component%3btwo, http://example.com/component%3dequals*

While delimiters can be used in an escaped and unescaped form in URIs with generally distinct meanings, basic CRIs (i.e., without percent-encoded text [Section 7.1](#)) only support one escapable delimiter character per component, which is the delimiter by which the component is split up in the CRI.

Note that the separators `.` (for authority parts), `/` (for paths), `&` (for query parameters) are special in that they are syntactic delimiters of their respective components in CRIs. Thus, the following examples are convertible to basic CRIs:

`https://interior%2edot/`

`https://example.com/path%2fcomponent/second-component`

`https://example.com/x?ampersand=%26&questionmark=?`

*https://alice@example.com/

The user information can be expressed in CRIs if the "userinfo" feature is present. The URI https://example.com is represented as [-4, [false, "", "example", "com"]]; the false serves as a marker that the next element is the userinfo.

The rules do not cater for unencoded ":" in userinfo, which is commonly considered a deprecated inclusion of a literal password.

Appendix C. Change Log

This section is to be removed before publishing as an RFC.

Changes from -08 to -09

- *Identify more esoteric features with a CDDL ".feature".
- *Clarify that well-formedness requires removing trailing nulls.
- *Fragments can contain PET.
- *Percent-encoded text in PET is treated as byte strings.
- *URIs with an authority but a completely empty path (e.g., http://example.com): CRIs with an authority component no longer always produce at least a slash in the path component.

For generic schemes, the conversion of scheme://example.com to a CRI is now possible because CRI produces a URI with an authority not followed by a slash following the updated rules of [Section 6.1](#). Schemes like http and coap do not distinguish between the empty path and the path containing a single slash when an authority is set (as recommended in [\[RFC3986\]](#)). For these schemes, that equivalence allows implementations to convert the just-a-slash URI to a CRI with a zero length path array (which, however, when converted back, does not produce a slash after the authority).

(Add an appendix "the small print" for more detailed discussion of pesky corner cases like this.)

Changes from -07 to -08

- *Fix the encoding of NOAUTH-NOSLASH / NOAUTH-LEADINGSLASH
- *Add URN and DID schemes, add example.

- *Add PET

- *Remove hopeless attempt to encode "remote trailing nulls" rule in CDDL (which is not a transformation language).

Changes from -06 to -07

- *More explicitly discuss constraints ([Section 2](#)), add examples ([Appendix B, Paragraph 6, Item 1](#)).

- *Make CDDL more explicit about special simple values.

- *Lots of gratuitous changes from XML2RFC redefinition of <tt> semantics.

Changes from -05 to -06

- *rework authority:

 - split reg-names at dots;

 - add optional zone identifiers [[RFC6874](#)] to IP addresses

Changes from -04 to -05

- *Simplify CBOR structure.

- *Add implementation status section.

Changes from -03 to -04:

- *Minor editorial improvements.

- *Renamed path.type/path-type to discard.

- *Renamed option to section, substructured into items.

- *Simplified the table "resolution-variables".

- *Use the CBOR structure inspired by Jim Schaad's proposals.

Changes from -02 to -03:

- *Expanded the set of supported schemes (#3).

- *Specified creation, normalization and comparison (#9).

- *Clarified the default value of the path.type option (#33).

- *Removed the append-relation path.type option (#41).

*Renumbered the remaining path.types.

*Renumbered the option numbers.

*Restructured the document.

*Minor editorial improvements.

Changes from -01 to -02:

*Changed the syntax of schemes to exclude upper case characters (#13).

*Minor editorial improvements (#34 #37).

Changes from -00 to -01:

*None.

Acknowledgements

CRIs were developed by Klaus Hartke for use in the Constrained RESTful Application Language (CoRAL). The current author team is completing this work with a view to achieve good integration with the potential use cases, both inside and outside of CoRAL.

Thanks to Christian Amsüss, Thomas Fossati, Ari Keränen, Jim Schaad, Dave Thaler and Marco Tiloca for helpful comments and discussions that have shaped the document.

Contributors

Klaus Hartke
Ericsson
Torshamnsgatan 23
SE-16483 Stockholm
Sweden

Email: klaus.hartke@ericsson.com

Authors' Addresses

Carsten Bormann (editor)
Universität Bremen TZI
Postfach 330440
D-28359 Bremen
Germany

Phone: [+49-421-218-63921](tel:+49-421-218-63921)

Email: cabo@tzi.org

Henk Birkholz
Fraunhofer SIT
Rheinstrasse 75
64295 Darmstadt
Germany

Email: henk.birkholz@sit.fraunhofer.de