

Dispatch
Internet-Draft
Updates: RFC6442 (if approved)
Intended status: Standards Track
Expires: August 14, 2017

R. Jesske
Deutsche Telekom
February 10, 2017

Location Parameter for the SIP Reason Header Field
draft-jesske-dispatch-reason-loc-q850-00.txt

Abstract

The SIP Reason header field is defined for Q.850 cause values. Some services in SIP networks may need to know the location where the call was released in the PSTN network to correctly interpret the reason of release.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 14, 2017.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Terminology	2
3. Rationale	2
4. Mechanism	3
5. Example	3
6. Privacy Considerations	4
7. Security Considerations	4
8. IANA Considerations	4
8.1. Registration of resp-location Parameter for reason header field	4
9. Acknowledgments	5
10. References	5
10.1. Normative References	5
10.2. Informative References	5
Author's Address	5

1. Introduction

The SIP Reason header field specification [RFC3326] describes a SIP header field that is used to indicate that a SIP request or response is carrying the reason of release. The location information identifies the part of the ISUP network where the call was released.

This document adds a location value parameter to the reason-extension parameter in [RFC3326] so that the location value can be interworked from the PSTN. The interworking from PTSN needs only to include the location received by the interworking gateway.

2. Terminology

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

3. Rationale

The primary intent of the parameter defined in this specification is for use in IMS networks defined by 3GPP but also open to be used by any other network. The purpose of this parameter is to transport the location of call release from the originating PSTN entity to the SIP entity receiving the response or BYE message containing the location of the call release. The ISDN location is defined in [Q.850].

4. Mechanism

As defined by [RFC3326] a Reason header field MAY appear in any request in a dialog, in any CANCEL request and in any response whose status code explicitly allows the presence of this header field. The syntax of the header field follows the standard SIP parameter syntax.

The mechanism employed adds a parameter with the location value defined in [Q.850] to the Reason header field that identifies the location of the call release in ISUP as defined in [Q.850]. The Augmented BNF (ABNF) [RFC5234] for this parameter is shown in Figure 1.

```
reason-extension  =/ resp-location
resp-location     = "location" EQUAL string
```

The values to be used as location are:

U	for user
LPN	for private network serving the local user
LN	for public network serving the local user
TN	for transit network
RLN	for public network serving the remote user
RPN	for private network serving the remote user
INTL	for international network
BI	for network beyond interworking point

Figure 1: resp-location

Depending on the direction the UAC or UAS shall include the resp-location when setting up the Reason header field with a Q.850 cause. This approach is only valid in cases when the ISUP Q.850 location is available.

5. Example

The following example shows a SIP 404 response message containing a Reason header field with a Q.850 cause value and a resp-location value. The 404 Response will be set up when a gateway receives an ISUP Release with a Q.850 cause set to 1 meaning "Unallocated (unassigned) number", i.e. the number is not known in the PSTN.

404 Not Found

SIP/2.0 404 Not Found

```
From: Alice <sips:alice@atlanta.example.com>;tag=1234567
To: Bob <sips:bob@biloxi.example.com>;tag=765432
Call-ID: 12345600@atlanta.example.com
CSeq: 1 INVITE
Reason: Q.850;cause=1;text="Unallocated (unassigned) number"; location=LN
Content-Length: 0
```

Figure 2: Example Location in Reason header field.

6. Privacy Considerations

This document doesn't change any of the privacy considerations described in [RFC3326]. While the addition of the resp-location parameter does provide an indicator of the entity that added the location in the signaling path this provides little more exposure than the Q.850 cause itself.

7. Security Considerations

This document doesn't change any of the security considerations described in [RFC3326]. The addition of the resp-location parameter does provide an indicator of the location where the call was release within the PSTN. This information may be used for specific location driven services but does not create any additional security constraints. But since the location is very imprecise the location value itself will not add any major security constrain. The use of this parameter is not restricted to a specific architecture.

8. IANA Considerations

8.1. Registration of resp-location Parameter for reason header field

This document calls for IANA to register a new SIP header parameter as per the guidelines in [RFC3261], which will be added to header sub-registry under <http://www.iana.org/assignments/sip-parameters>.

Header Field: Reason

Parameter Name: resp-location

9. Acknowledgments

Thanks to Michael Kreipel for his comments and review.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <<http://www.rfc-editor.org/info/rfc3261>>.
- [RFC3326] Schulzrinne, H., Oran, D., and G. Camarillo, "The Reason Header Field for the Session Initiation Protocol (SIP)", RFC 3326, DOI 10.17487/RFC3326, December 2002, <<http://www.rfc-editor.org/info/rfc3326>>.
- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, DOI 10.17487/RFC5234, January 2008, <<http://www.rfc-editor.org/info/rfc5234>>.

10.2. Informative References

- [Q.850] INTERNATIONAL TELECOMMUNICATION UNION, "Usage of cause and location in the Digital Subscriber Signalling System No. 1 and the Signalling System No. 7 ISDN User Part", Q 850, May 1998.

Author's Address

Roland Jesske
Deutsche Telekom
Heinrich-Hertz Str, 3-7
Darmstadt 64295
Germany

Email: r.jesske@telekom.de
URI: www.telekom.de

Network Working Group
Internet-Draft
Obsoletes: 5988 (if approved)
Intended status: Standards Track
Expires: August 7, 2017

M. Nottingham
February 3, 2017

Web Linking
draft-nottingham-rfc5988bis-04

Abstract

This specification defines a way to indicate the relationships between resources on the Web ("links") and the type of those relationships ("link relation types").

It also defines the serialisation of such links in HTTP headers with the Link header field.

Note to Readers

This is a work-in-progress to revise RFC5988.

The issues list can be found at <https://github.com/mnot/I-D/labels/rfc5988bis> .

The most recent (often, unpublished) draft is at <https://mnot.github.io/I-D/rfc5988bis/> .

Recent changes are listed at <https://github.com/mnot/I-D/commits/gh-pages/rfc5988bis> .

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 7, 2017.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	3
2.	Notational Conventions	3
3.	Links	4
4.	Link Relation Types	5
4.1.	Registered Relation Types	5
4.1.1.	Registering Link Relation Types	6
4.1.2.	Registration Request Processing	7
4.2.	Extension Relation Types	7
5.	Target Attributes	8
6.	Link Serialisation in HTTP Headers	8
6.1.	Link Target	9
6.2.	Link Context	9
6.3.	Relation Type	9
6.4.	Target Attributes	10
6.4.1.	Serialisation-Defined Attributes	10
6.4.2.	Extension Attributes	11
6.5.	Examples	12
7.	IANA Considerations	13
7.1.	Link HTTP Header Field Registration	13

7.2. Link Relation Type Registry	13
8. Security Considerations	14
9. Internationalisation Considerations	14
10. References	14
10.1. Normative References	15
10.2. Informative References	16
Appendix A. Notes on Other Link Serialisations	17
A.1. Link Serialisation in HTML	17
A.2. Link Serialisation in Atom	17
Appendix B. Algorithm for Parsing Link Headers	18
Appendix C. Changes from RFC5988	21
Author's Address	22

1. Introduction

This specification defines a way to indicate the relationships between resources on the Web ("links") and the type of those relationships ("link relation types").

HTML [W3C.REC-html5-20141028] and Atom [RFC4287] both have well-defined concepts of linking; this specification generalises this into a framework that encompasses linking in these formats and (potentially) elsewhere.

Furthermore, this specification formalises an HTTP header field for conveying such links, having been originally defined in Section 19.6.2.4 of [RFC2068], but removed from [RFC2616].

2. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14, [RFC2119], as scoped to those conformance targets.

This document uses the Augmented Backus-Naur Form (ABNF) notation of [RFC7230], including the #rule, and explicitly includes the following rules from it: quoted-string, token, SP (space), BWS (bad whitespace), OWS (optional whitespace), RWS (required whitespace) LOALPHA, DIGIT.

Additionally, the following rules are included from [RFC3986]: URI and URI-Reference; from [RFC6838]: type-name and subtype-name; from [W3C.CR-css3-mediaqueries-20090915]: media_query_list; and from [RFC5646]: Language-Tag..

3. Links

In this specification, a link is a typed connection between two resources, and is comprised of:

- o A `_link context_`,
- o a `_link relation type_` (Section 4),
- o a `_link target_`, and
- o optionally, `_target attributes_` (Section 5).

A link can be viewed as a statement of the form "{link context} has a {link relation type} resource at {link target}, which has {target attributes}".

Link contexts and link targets are both IRIs [RFC3987]. However, in the common case, the link context will also be a URI [RFC3986], because many protocols (such as HTTP) do not support dereferencing IRIs. Likewise, the link target will sometimes be converted to a URI (see [RFC3987], Section 3.1) in places that do not support IRIs (such as the Link header field defined in Section 6).

This specification does not place restrictions on the cardinality of links; there can be multiple links to and from a particular target, and multiple links of the same or different types between a given context and target. Likewise, the relative ordering of links in any particular serialisation, or between serialisations (e.g., the Link header field and in-content links) is not specified or significant in this specification; applications that wish to consider ordering significant can do so.

Links are conveyed in `_link serialisations_`; they are the "bytes on the wire", and can occur in various forms. For example, Atom [RFC4287] and HTML [W3C.REC-html5-20141028] both defined serialisations of links into their respective formats, and Section 6 defines how to serialise links in HTTP header fields.

This specification does not define a general syntax for links across different serialisations, nor does it mandate a specific context for any given link; it is expected that serialisations of links will specify both aspects.

Finally, links are consumed by `_link applications_`. Generally, an application will define the link relation types it uses, along with the serialisations that they might occur within. For example, the application "Web browsing" looks for the "stylesheet" link relation

type in the HTML link serialisation, whereas the application "AtomPub" uses the "edit" and "edit-media" link relations.

4. Link Relation Types

In the simplest case, a link relation type identifies the semantics of a link. For example, a link with the relation type "copyright" indicates that the resource identified by the link target is a statement of the copyright terms applying to the current link context.

Link relation types can also be used to indicate that the target resource has particular attributes, or exhibits particular behaviours; for example, a "service" link implies that the identified resource is part of a defined protocol (in this case, a service description).

Relation types are not to be confused with media types [RFC6838]; they do not identify the format of the representation that results when the link is dereferenced. Rather, they only describe how the current context is related to another resource.

Relation types SHOULD NOT infer any additional semantics based upon the presence or absence of another link relation type, or its own cardinality of occurrence. An exception to this is the combination of the "alternate" and "stylesheet" registered relation types, which has special meaning in HTML for historical reasons.

There are two kinds of relation types: registered and extension.

4.1. Registered Relation Types

Well-defined relation types can be registered as tokens for convenience and/or to promote reuse by other applications, using the procedure in Section 4.1.1.

Registered relation type names MUST conform to the reg-rel-type rule (see Section 6.3), and MUST be compared character-by-character in a case-insensitive fashion. They SHOULD be appropriate to the specificity of the relation type; i.e., if the semantics are highly specific to a particular application, the name should reflect that, so that more general names are available for less specific use.

Registered relation types MUST NOT constrain the media type of the link context, and MUST NOT constrain the available representation media types of the link target. However, they can specify the behaviours and properties of the target resource (e.g., allowable

HTTP methods, request and response media types that must be supported).

Historically, registered relation types have been identified with a URI [RFC3986] by prefixing their names with an application-defined base URI (e.g., see Appendix A.2). This practice is NOT RECOMMENDED, because the resulting strings will not be considered equivalent to the registered relation types by other processors. Applications that do use such URIs internally MUST NOT use them in link serialisations that do not explicitly accommodate them.

4.1.1. Registering Link Relation Types

Any party can request registration of a link relation type.

Registration requests can be sent to the "link-relations@ietf.org" mailing list. The Expert(s) MAY establish alternate means of requesting registrations, which SHOULD be linked to from the registry page.

Registration requests consist of at least the following information:

- o ***Relation Name***: The name of the relation type
- o ***Description***: A short English description of the type's semantics. It SHOULD be stated in terms of the relationship between the link context and link target.
- o ***Reference***: Reference to the document that specifies the link relation type, preferably including a URI that can be used to retrieve a copy of the document. An indication of the relevant section(s) MAY also be included, but is not required.

The Expert(s) MAY define additional fields to be collected in the registry.

General requirements for registered relation types are described in Section 4.1.

Registrations MUST reference a freely available, stable specification.

Note that relation types can be registered by third parties (including the Expert(s)), if the Expert(s) determine that an unregistered relation type is widely deployed and not likely to be registered in a timely manner.

4.1.2. Registration Request Processing

Relation types are registered on the advice of a Designated Expert (appointed by the IESG or their delegate), with a Specification Required (using terminology from [RFC5226]).

The goal of the registry is to reflect common use of links on the Internet. Therefore, the Expert(s) SHOULD be strongly biased towards approving registrations, unless they are abusive, frivolous, not likely to be used on the Internet, or actively harmful to the Internet and/or the Web (not merely aesthetically displeasing, or architecturally dubious).

The Expert(s) MUST clearly identify any issues which cause a registration to be refused. Advice about the syntax or semantics of a proposed link relation type can be given, but if it does not block registration, this SHOULD be explicitly stated.

When a request is approved, the Expert(s) will inform IANA, and the registration will be processed. The IESG is the final arbiter of any objection.

4.2. Extension Relation Types

Applications that don't wish to register a relation type can use an extension relation type, which is a URI [RFC3986] that uniquely identifies the relation type. Although the URI can point to a resource that contains a definition of the semantics of the relation type, clients SHOULD NOT automatically access that resource to avoid overburdening its server.

The URI used for an extension relation type SHOULD be under the control of the person or party defining it, or be delegated to them.

When extension relation types are compared, they MUST be compared as strings (after converting to URIs if serialised in a different format) in a case-insensitive fashion, character-by-character. Because of this, all-lowercase URIs SHOULD be used for extension relations.

Note that while extension relation types are required to be URIs, a serialisation of links can specify that they are expressed in another form, as long as they can be converted to URIs.

5. Target Attributes

Target attributes are a list of key/value pairs that describe the link or its target; for example, a media type hint.

They can be defined both by individual link relation types and by link serialisations.

This specification does not attempt to coordinate the name of target attributes, their cardinality or use. Serialisations SHOULD coordinate their target attributes to avoid conflicts in semantics or syntax.

The names of target attributes SHOULD conform to the token rule, but SHOULD NOT include any of the characters "%", "'" or "*", for portability across serializations, and MUST be compared in a case-insensitive fashion.

Target attribute definitions SHOULD specify:

- o The serialisation of their values into Unicode or a subset thereof, to maximise their chances of portability across link serialisations.
- o The semantics and error handling of multiple occurrences of the target attribute on a given link.

This specification does define target attributes for use in the Link HTTP header field in Section 6.4.

6. Link Serialisation in HTTP Headers

The Link header field provides a means for serialising one or more links into HTTP headers.

The ABNF for the field value is given below:

```
Link          = #link-value
link-value   = "<" URI-Reference ">" *( OWS ";" OWS link-param )
link-param   = token BWS "=" BWS ( token / quoted-string )
```

Note that any "link-param" can be generated with values using either the "token" or the "quoted-string" syntax, and therefore recipients MUST be able to parse both forms. Individual "link-param"s specify their syntax in terms of the value after any necessary unquoting (as per [RFC7230], Section 3.2.6).

This specification defines the link-params "rel", "anchor", "rev", "hreflang", "media", "title", "title*", and "type"; see Section 6.2, Section 6.3 and Section 6.4.

6.1. Link Target

Each link-value conveys one target IRI as a URI-Reference (after conversion to one, if necessary; see [RFC3987], Section 3.1) inside angle brackets ("`<`"). If the URI-Reference is relative, parsers MUST resolve it as per [RFC3986], Section 5. Note that any base IRI from the message's content is not applied.

6.2. Link Context

By default, the context of a link conveyed in the Link header field is identity of the representation it is associated with, as defined in [RFC7231], Section 3.1.4.1, serialised as a URI.

When present, the anchor parameter overrides this with another URI, such as a fragment of this resource, or a third resource (i.e., when the anchor value is an absolute URI). If the anchor parameter's value is a relative URI, parsers MUST resolve it as per [RFC3986], Section 5. Note that any base URI from the body's content is not applied.

The ABNF for the "anchor" parameter's value is:

```
URI-Reference
```

Consuming implementations can choose to ignore links with an anchor parameter. For example, the application in use might not allow the link context to be assigned to a different resource. In such cases, the entire link is to be ignored; consuming implementations MUST NOT process the link without applying the anchor.

Note that depending on HTTP status code and response headers, the link context might be "anonymous" (i.e., no link context is available). For instance, this is the case on a 404 response to a GET request.

6.3. Relation Type

The relation type of a link conveyed in the Link header field is conveyed in the "rel" parameter's value. The "rel" parameter MUST NOT appear more than once in a given link-value; occurrences after the first MUST be ignored by parsers.

The "rev" parameter has been used in the past to indicate that the semantics of the relationship are in the reverse direction. That is, a link from A to B with REL="X" expresses the same relationship as a link from B to A with REV="X". "rev" is deprecated by this specification because it often confuses authors and readers; in most cases, using a separate relation type is preferable.

The ABNF for the "rel" and "rev" parameters' values is:

```
relation-type *( 1*SP relation-type )
```

where:

```
relation-type = reg-rel-type | ext-rel-type
reg-rel-type  = LOALPHA *( LOALPHA | DIGIT | "." | "-" )
ext-rel-type  = URI
```

Note that extension relation types are REQUIRED to be absolute URIs in Link headers, and MUST be quoted if they contain a semicolon (";") or comma (",") (as these characters are used as delimiters in the header field itself).

6.4. Target Attributes

The Link header field defines several target attributes specific to this serialisation, and also allows extension target attributes. Target attributes are serialised in the Link header field as parameters (see [RFC7231], Section 3.1.1.1 for the definition of their syntax).

6.4.1. Serialisation-Defined Attributes

The "hreflang", "media", "title", "title*", and "type" link-params can be translated to serialisation-defined target attributes for the link.

The "hreflang" attribute, when present, is a hint indicating what the language of the result of dereferencing the link should be. Note that this is only a hint; for example, it does not override the Content-Language header field of a HTTP response obtained by actually following the link. Multiple "hreflang" attributes on a single link-value indicate that multiple languages are available from the indicated resource.

The ABNF for the "hreflang" parameter's value is:

```
Language-Tag
```

The "media" attribute, when present, is used to indicate intended destination medium or media for style information (see [W3C.REC-html5-20141028], Section 4.2.4). Its value MUST be quoted if it contains a semicolon (";") or comma (","). There MUST NOT be more than one "media" attribute in a link-value; occurrences after the first MUST be ignored by parsers.

The ABNF for the "media" parameter's value is:

```
media_query_list
```

The "title" attribute, when present, is used to label the destination of a link such that it can be used as a human-readable identifier (e.g., a menu entry) in the language indicated by the Content-Language header field (if present). The "title" attribute MUST NOT appear more than once in a given link; occurrences after the first MUST be ignored by parsers.

The "title*" link-param can be used to encode this attribute in a different character set, and/or contain language information as per [I-D.ietf-httpbis-rfc5987bis]. The "title*" link-param MUST NOT appear more than once in a given link-value; occurrences after the first MUST be ignored by parsers. If the attribute does not contain language information, its language is indicated by the Content-Language header field (when present).

If both the "title" and "title*" link-param appear in a link, processors SHOULD use the "title*" link-param's value for the "title" attribute.

The "type" attribute, when present, is a hint indicating what the media type of the result of dereferencing the link should be. Note that this is only a hint; for example, it does not override the Content-Type header field of a HTTP response obtained by actually following the link. The "type" attribute MUST NOT appear more than once in a given link-value; occurrences after the first MUST be ignored by parsers.

The ABNF for the "type" parameter's value is:

```
type-name "/" subtype-name
```

6.4.2. Extension Attributes

Other link-params are link-extensions, and are to be considered as target attributes.

Such target attributes MAY be defined to use the encoding in [I-D.ietf-httpbis-rfc5987bis] (e.g., "example" and "example_"). When both forms are present, they SHOULD be considered to be the same target attribute; processors SHOULD use the value of the name ending in "_" (after [I-D.ietf-httpbis-rfc5987bis] decoding), but MAY fall back to the other value if there is an error in decoding it, or if they do not support decoding.

6.5. Examples

For example:

```
Link: <http://example.com/TheBook/chapter2>; rel="previous";  
      title="previous chapter"
```

indicates that "chapter2" is previous to this resource in a logical navigation path.

Similarly,

```
Link: </>; rel="http://example.net/foo"
```

indicates that the root resource ("/") is related to this resource with the extension relation type "http://example.net/foo".

The example below shows an instance of the Link header field encoding multiple links, and also the use of RFC 5987 encoding to encode both non-ASCII characters and language information.

```
Link: </TheBook/chapter2>;  
      rel="previous"; title*=UTF-8'de'letztes%20Kapitel,  
      </TheBook/chapter4>;  
      rel="next"; title*=UTF-8'de'n%c3%a4chstes%20Kapitel
```

Here, both links have titles encoded in UTF-8, use the German language ("de"), and the second link contains the Unicode code point U+00E4 ("LATIN SMALL LETTER A WITH DIAERESIS").

Note that link-values can convey multiple links between the same link target and link context; for example:

```
Link: <http://example.org/>;  
      rel="start http://example.net/relation/other"
```

Here, the link to "http://example.org/" has the registered relation type "start" and the extension relation type "http://example.net/relation/other".

7. IANA Considerations

In addition to the actions below, IANA should terminate the Link Relation Application Data Registry, as it has not been used, and future use is not anticipated.

7.1. Link HTTP Header Field Registration

This specification updates the Message Header registry entry for "Link" in HTTP [RFC3864] to refer to this document.

Header field: Link
Applicable protocol: http
Status: standard
Author/change controller:
 IETF (iesg@ietf.org)
 Internet Engineering Task Force
Specification document(s):
 [RFC&rfc.number;]

7.2. Link Relation Type Registry

This specification updates the registration procedures for the Link Relation Type registry; see Section 4.1.1. The Expert(s) and IANA will interact as outlined below.

IANA will direct any incoming requests regarding the registry to this document and, if defined, the processes established by the Expert(s); typically, this will mean referring them to the registry Web page.

The Expert(s) will provide registry data to IANA in an agreed form (e.g. a specific XML format). IANA will publish:

- o The raw registry data
- o The registry data, transformed into HTML
- o The registry data in any alternative formats provided by the Expert(s)

Each published document will be at a URL agreed to by IANA and the Expert(s), and IANA will set HTTP response headers on them as (reasonably) requested by the Expert(s).

Additionally, the HTML generated by IANA will:

- o Take directions from the Expert(s) as to the content of the HTML page's introductory text

- o Include a stable HTML fragment identifier for each registered link relation

All registry data documents MUST include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions (<http://trustee.ietf.org/license-info>).

8. Security Considerations

The content of the Link header field is not secure, private or integrity-guaranteed, and due caution should be exercised when using it. Use of Transport Layer Security (TLS) with HTTP ([RFC2818] and [RFC2817]) is currently the only end-to-end way to provide such protection.

Link applications ought to consider the attack vectors opened by automatically following, trusting, or otherwise using links gathered from HTTP headers. In particular, Link header fields that use the "anchor" parameter to associate a link's context with another resource should be treated with due caution.

The Link header field makes extensive use of IRIs and URIs. See [RFC3987] for security considerations relating to IRIs. See [RFC3986] for security considerations relating to URIs. See [RFC7230] for security considerations relating to HTTP headers.

9. Internationalisation Considerations

Link targets may need to be converted to URIs in order to express them in serialisations that do not support IRIs. This includes the Link HTTP header field.

Similarly, the anchor parameter of the Link header field does not support IRIs, and therefore IRIs must be converted to URIs before inclusion there.

Relation types are defined as URIs, not IRIs, to aid in their comparison. It is not expected that they will be displayed to end users.

Note that registered Relation Names are required to be lower-case ASCII letters.

10. References

10.1. Normative References

- [I-D.ietf-httpbis-rfc5987bis]
Reschke, J., "Indicating Character Encoding and Language for HTTP Header Field Parameters", draft-ietf-httpbis-rfc5987bis-04 (work in progress), November 2016.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3864] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", BCP 90, RFC 3864, DOI 10.17487/RFC3864, September 2004, <<http://www.rfc-editor.org/info/rfc3864>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<http://www.rfc-editor.org/info/rfc3986>>.
- [RFC3987] Duerst, M. and M. Suignard, "Internationalized Resource Identifiers (IRIs)", RFC 3987, DOI 10.17487/RFC3987, January 2005, <<http://www.rfc-editor.org/info/rfc3987>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC5646] Phillips, A., Ed. and M. Davis, Ed., "Tags for Identifying Languages", BCP 47, RFC 5646, DOI 10.17487/RFC5646, September 2009, <<http://www.rfc-editor.org/info/rfc5646>>.
- [RFC6838] Freed, N., Klensin, J., and T. Hansen, "Media Type Specifications and Registration Procedures", BCP 13, RFC 6838, DOI 10.17487/RFC6838, January 2013, <<http://www.rfc-editor.org/info/rfc6838>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<http://www.rfc-editor.org/info/rfc7230>>.

- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, DOI 10.17487/RFC7231, June 2014, <<http://www.rfc-editor.org/info/rfc7231>>.
- [W3C.CR-css3-mediaqueries-20090915]
Lie, H., Celik, T., Glazman, D., and A. Kesteren, "Media Queries", World Wide Web Consortium CR CR-css3-mediaqueries-20090915, September 2009, <<http://www.w3.org/TR/2009/CR-css3-mediaqueries-20090915>>.

10.2. Informative References

- [RFC2068] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2068, DOI 10.17487/RFC2068, January 1997, <<http://www.rfc-editor.org/info/rfc2068>>.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, DOI 10.17487/RFC2616, June 1999, <<http://www.rfc-editor.org/info/rfc2616>>.
- [RFC2817] Khare, R. and S. Lawrence, "Upgrading to TLS Within HTTP/1.1", RFC 2817, DOI 10.17487/RFC2817, May 2000, <<http://www.rfc-editor.org/info/rfc2817>>.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, DOI 10.17487/RFC2818, May 2000, <<http://www.rfc-editor.org/info/rfc2818>>.
- [RFC4287] Nottingham, M., Ed. and R. Sayre, Ed., "The Atom Syndication Format", RFC 4287, DOI 10.17487/RFC4287, December 2005, <<http://www.rfc-editor.org/info/rfc4287>>.
- [W3C.REC-html-rdfa-20150317]
Sporny, M., "HTML+RDFa 1.1 - Second Edition", World Wide Web Consortium Recommendation REC-html-rdfa-20150317, March 2015, <<http://www.w3.org/TR/2015/REC-html-rdfa-20150317>>.
- [W3C.REC-html5-20141028]
Hickson, I., Berjon, R., Faulkner, S., Leithead, T., Navara, E., O'Connor, T., and S. Pfeiffer, "HTML5", World Wide Web Consortium Recommendation REC-html5-20141028, October 2014, <<http://www.w3.org/TR/2014/REC-html5-20141028>>.

Appendix A. Notes on Other Link Serialisations

Header fields (Section 6) are only one serialisation of links; other specifications have defined alternative serialisations.

A.1. Link Serialisation in HTML

HTML motivated the original syntax of the Link header field, and many of the design decisions in this document are driven by a desire to stay compatible with it.

In HTML, the link element can be mapped to links as specified here by using the "href" attribute for the target URI, and "rel" to convey the relation type, as in the Link header field. The context of the link is the URI associated with the entire HTML document. HTML also defines several attributes on links that can be seen as target attributes, including "media", "hreflang", "type" and "sizes".

HTML5 ([W3C.REC-html5-20141028]) Section 4.8 defines modern HTML links. That document links to the Microformats Wiki as a registry; over time, the IANA registry ought to mirror its contents, and ideally eventually replace it (although that depends on the HTML community).

Surveys of existing HTML content have shown that unregistered link relation types that are not URIs are (perhaps inevitably) common. Consuming HTML implementations ought not consider such unregistered short links to be errors, but rather relation types with a local scope (i.e., their meaning is specific and perhaps private to that document).

Finally, the HTML specification gives a special meaning when the "alternate" relation type coincides with other relation types in the same link. Such links ought to be serialised in the Link header field using a single list of relation-types (e.g., rel="alternate stylesheet") to preserve this relationship.

A.2. Link Serialisation in Atom

Atom [RFC4287] is a link serialisation that conveys links in the atom:link element, with the "href" attribute indicating the link target and the "rel" attribute containing the relation type. The context of the link is either a feed locator or an entry ID, depending on where it appears; generally, feed-level links are obvious candidates for transmission as a Link header field.

When serialising an atom:link into a Link header field, it is necessary to convert link targets (if used) to URIs.

Atom defines extension relation types in terms of IRIs. This specification re-defines them as URIs, to simplify and reduce errors in their comparison.

Atom allows registered link relation types to be serialised as absolute URIs using a prefix, "http://www.iana.org/assignments/relation/". This prefix is specific to the Atom serialisation.

Furthermore, link relation types are always compared in a case-sensitive fashion; therefore, registered link relation types SHOULD be converted to their registered form (usually, lowercase) when serialised in an Atom document.

Note also that while the Link header field allows multiple relations to be serialised in a single link, atom:link does not. In this case, a single link-value may map to several atom:link elements.

As with HTML, atom:link defines some attributes that are not explicitly mirrored in the Link header field syntax, but they can also be used as link-extensions to maintain fidelity.

Appendix B. Algorithm for Parsing Link Headers

Given a HTTP header field-value "field_value" as a string assuming ASCII encoding, the following algorithm can be used to parse it into the model described by this specification:

1. Let "links" be an empty list.
2. Create "link_strings" by splitting "field_value" on "," characters, excepting "," characters within quoted strings as per [RFC7230], Section 3.2.6, or which form part of link's URI-Reference (i.e. between "<" and ">" characters where the "<" is immediately preceded by OWS and either a "," character or the beginning of the "field_value" string).
3. For each "link_string" in "link_strings":
 1. Let "target_string" be the string between the first "<" and first ">" characters in "link_string". If they do not appear, or do not appear in that order, fail parsing.
 2. Let "rest" be the remaining characters (if any) after the first ">" character in "link_string".
 3. Split "rest" into an array of strings "parameter_strings", on the ";" character, excepting ";" characters within quoted strings as per [RFC7230], Section 3.2.6.

4. Let "link_parameters" be an empty array.
5. For each item "parameter" in "parameter_strings":
 1. Remove OWS from the beginning and end of "parameter".
 2. Skip this item if "parameter" matches the empty string ("").
 3. Split "parameter" into "param_name" and "param_value" on the first "=" character. If "parameter" does not contain "=", let "param_name" be "parameter" and "param_value" be null.
 4. Remove OWS from the end of "param_name" and the beginning of "param_value".
 5. Case-normalise "param_name" to lowercase.
 6. If the first and last characters of "param_value" are both DQUOTE:
 1. Remove the first and last characters of "param_value".
 2. Replace quoted-pairs within "param_value" with the octet following the backslash, as per [RFC7230], Section 3.2.6.
 7. If the last character of "param_name" is an asterisk ("*"), decode "param_value" according to [I-D.ietf-httpbis-rfc5987bis]. Skip this item if an unrecoverable error is encountered.
 8. Append the tuple ("param_name", "param_value") to "link_parameters".
6. Let "target" be the result of relatively resolving (as per [RFC3986], Section 5.2) "target_string". Note that any base URI carried in the payload body is NOT used.
7. Let "relations_string" be the second item of the first tuple of "link_parameters" whose first item matches the string "rel", or the empty string (") if it is not present.
8. Split "relations_string" into an array of strings "relation_types", on RWS (removing all whitespace in the process).

9. Let "context_string" be the second item of the first tuple of "link_parameters" whose first item matches the string "anchor". If it is not present, "context_string" is the identity of the representation carrying the Link header [RFC7231], Section 3.1.4.1, serialised as a URI. Where the identity is "anonymous" "context_string" is null.
10. Let "context" be the result of relatively resolving (as per [RFC3986], Section 5.2) "context_string", unless "context_string" is null in which case "context" is null. Note that any base URI carried in the payload body is NOT used.
11. Let "target_attributes" be an empty array.
12. For each tuple ("param_name", "param_value") of "link_parameters":
 1. If "param_name" matches "rel" or "anchor", skip this tuple.
 2. If "param_name" matches "media", "title", "title*" or "type" and "target_attributes" already contains a tuple whose first element matches the value of "param_name", skip this tuple.
 3. Append ("param_name", "param_value") to "target_attributes".
13. Let "star_param_names" be the set of "param_name"s in the ("param_name", "param_value") tuples of "link_parameters" where the last character of "param_name" is an asterisk ("*").
14. For each "star_param_name" in "star_param_names":
 1. Let "base_param_name" be "star_param_name" with the last character removed.
 2. If the implementation does not choose to support an internationalised form of a parameter named "base_param_name" for any reason (including, but not limited to, it being prohibited by the parameter's specification), remove all tuples from "link_parameters" whose first member is "star_param_name" and skip to the next "star_param_name".

3. Remove all tuples from "link_parameters" whose first member is "base_param_name".
 4. Change the first member of all tuples in "link_parameters" whose first member is "star_param_name" to "base_param_name".
15. For each "relation_type" in "relation_types":
1. Case-normalise "relation_type" to lowercase.
 2. Append a link object to "links" with the target "target", relation type of "relation_type", context of "context", and target attributes "target_attributes".
4. Return "links".

Appendix C. Changes from RFC5988

This specification has the following differences from its predecessor, RFC5988:

- o The initial relation type registrations were removed, since they've already been registered by 5988.
- o The introduction has been shortened.
- o The Link Relation Application Data Registry has been removed.
- o Incorporated errata.
- o Updated references.
- o Link cardinality was clarified.
- o Terminology was changed from "target IRI" and "context IRI" to "link target" and "link context" respectively.
- o Made assigning a URI to registered relation types serialisation-specific.
- o Removed misleading statement that the link header field is semantically equivalent to HTML and Atom links.
- o More carefully defined how the Experts and IANA should interact.
- o More carefully defined and used "link serialisations" and "link applications."

- o Clarified the cardinality of target attributes (generically and for "type").
- o Corrected the default link context for the Link header field, to be dependent upon the identity of the representation (as per RFC7231).
- o Defined a suggested parsing algorithm for the Link header.
- o The value space of target attributes and their definition has been specified.
- o The ABNF has been updated to be compatible with [RFC7230]. In particular, whitespace is now explicit.
- o Some parameters on the HTTP header field can now appear as a token.
- o Handling of quoted strings is now defined by [RFC7230].
- o The "type" header field parameter now needs to be quoted (as "token" does not allow "/").

Author's Address

Mark Nottingham

E-Mail: mnot@mnot.net

URI: <https://www.mnot.net/>

Network Working Group
Internet-Draft
Obsoletes: 5988 (if approved)
Intended status: Standards Track
Expires: February 12, 2018

M. Nottingham
August 11, 2017

Web Linking
draft-nottingham-rfc5988bis-08

Abstract

This specification defines a model for the relationships between resources on the Web ("links") and the type of those relationships ("link relation types").

It also defines the serialisation of such links in HTTP headers with the Link header field.

Note to Readers

RFC EDITOR: please remove this section before publication

This is a work-in-progress to revise RFC5988.

The issues list can be found at <https://github.com/mnot/I-D/labels/rfc5988bis>.

The most recent (often, unpublished) draft is at <https://mnot.github.io/I-D/rfc5988bis/>.

Recent changes are listed at <https://github.com/mnot/I-D/commits/gh-pages/rfc5988bis>.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on February 12, 2018.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	3
1.1.	Notational Conventions	3
1.2.	Conformance and Error Handling	4
2.	Links	4
2.1.	Link Relation Types	5
2.1.1.	Registered Relation Types	5
2.1.2.	Extension Relation Types	7
2.2.	Target Attributes	8
3.	Link Serialisation in HTTP Headers	8
3.1.	Link Target	9
3.2.	Link Context	9
3.3.	Relation Type	10
3.4.	Target Attributes	11
3.4.1.	Serialisation-Defined Attributes	11
3.4.2.	Extension Attributes	12
3.5.	Link Header Field Examples	12
4.	IANA Considerations	14

4.1. Link HTTP Header Field Registration	14
4.2. Link Relation Type Registry	14
4.3. Link Relation Application Data Registry	14
5. Security Considerations	14
6. Internationalisation Considerations	15
7. References	15
7.1. Normative References	15
7.2. Informative References	17
Appendix A. Notes on Other Link Serialisations	18
A.1. Link Serialisation in HTML	18
A.2. Link Serialisation in Atom	18
Appendix B. Algorithms for Parsing Link Header Fields	19
B.1. Parsing a Header Set for Links	19
B.2. Parsing a Link Field Value	20
B.3. Parsing Parameters	22
B.4. Parsing a Quoted String	23
Appendix C. Changes from RFC5988	24
Author's Address	25

1. Introduction

This specification defines a model for the relationships between resources on the Web ("links") and the type of those relationships ("link relation types").

HTML [W3C.REC-html5-20141028] and Atom [RFC4287] both have well-defined concepts of linking; Section 2 generalises this into a framework that encompasses linking in these formats and (potentially) elsewhere.

Furthermore, Section 3 defines an HTTP header field for conveying such links.

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.

This document uses the Augmented Backus-Naur Form (ABNF) notation of [RFC7230], including the #rule, and explicitly includes the following rules from it: quoted-string, token, SP (space), BWS (bad whitespace), OWS (optional whitespace), RWS (required whitespace) LOALPHA, DIGIT.

Additionally, the following rules are included from [RFC3986]: URI and URI-Reference; from [RFC6838]: type-name and subtype-name; from [W3C.REC-css3-mediaqueries-20120619]: media-query-list; and from [RFC5646]: Language-Tag.

1.2. Conformance and Error Handling

The requirements regarding conformance and error handling highlighted in [RFC7230], Section 2.5 apply to this document.

2. Links

In this specification, a link is a typed connection between two resources, and is comprised of:

- o A `_link context_`,
- o a `_link relation type_` (Section 2.1),
- o a `_link target_`, and
- o optionally, `_target attributes_` (Section 2.2).

A link can be viewed as a statement of the form "`_link context_` has a `_link relation type_` resource at `_link target_`, which has `_target attributes_`".

For example, `https://www.example.com/` has a "canonical" resource at `https://example.com`, which has a "type" of "text/html".

Link contexts and link targets are both IRIs [RFC3987]. However, in the common case, the link context will also be a URI [RFC3986], because many protocols (such as HTTP) do not support dereferencing IRIs. Likewise, the link target will be sometimes be converted to a URI (see [RFC3987], Section 3.1) in serialisations that do not support IRIs (such as the Link header field defined in Section 3).

This specification does not place restrictions on the cardinality of links; there can be multiple links to and from a particular target, and multiple links of the same or different types between a given context and target. Likewise, the relative ordering of links in any particular serialisation, or between serialisations (e.g., the Link header field and in-content links) is not specified or significant in this specification; applications that wish to consider ordering significant can do so.

Links are conveyed in `_link serialisations_`; they are the "bytes on the wire", and can occur in various forms. For example, Atom

[RFC4287] and HTML [W3C.REC-html5-20141028] both defined serialisations of links into their respective formats, and Section 3 defines how to serialise links in HTTP header fields.

This specification does not define a general syntax for links across different serialisations, nor does it mandate a specific context for any given link; it is expected that serialisations of links will specify both aspects.

Finally, links are used by `_link applications_`. Generally, an application will define the link relation type(s) it uses, along with the serialisation(s) that they might occur within. For example, the application "Web browsing" looks for the "stylesheet" link relation type in the HTML link serialisation (and optionally in the Link header field), whereas the application "AtomPub" uses the "edit" and "edit-media" link relations in the Atom serialisation.

2.1. Link Relation Types

In the simplest case, a link relation type identifies the semantics of a link. For example, a link with the relation type "copyright" indicates that the current link context has a copyright resource at the link target.

Link relation types can also be used to indicate that the target resource has particular attributes, or exhibits particular behaviours; for example, a "service" link implies that the link target can be used as part of a defined protocol (in this case, a service description).

Relation types are not to be confused with media types [RFC2046]; they do not identify the format of the representation that results when the link is dereferenced. Rather, they only describe how the current context is related to another resource.

Relation types SHOULD NOT infer any additional semantics based upon the presence or absence of another link relation type, or its own cardinality of occurrence. An exception to this is the combination of the "alternate" and "stylesheet" registered relation types, which has special meaning in HTML for historical reasons.

There are two kinds of relation types: registered and extension.

2.1.1. Registered Relation Types

Well-defined relation types can be registered as tokens for convenience and/or to promote reuse by other applications, using the procedure in Section 2.1.1.1.

Registered relation type names MUST conform to the reg-rel-type rule (see Section 3.3), and MUST be compared character-by-character in a case-insensitive fashion. They SHOULD be appropriate to the specificity of the relation type; i.e., if the semantics are highly specific to a particular application, the name should reflect that, so that more general names are available for less specific use.

Registered relation types MUST NOT constrain the media type of the link context, and MUST NOT constrain the available representation media types of the link target. However, they can specify the behaviours and properties of the target resource (e.g., allowable HTTP methods, request and response media types that are required be supported).

Historically, registered relation types have been identified with a URI [RFC3986] by prefixing their names with an application-defined base URI (e.g., see Appendix A.2). This practice is NOT RECOMMENDED, because the resulting strings will not be considered equivalent to the registered relation types by other applications. Applications that do use such URIs internally MUST NOT use them in link serialisations that do not explicitly accommodate them.

2.1.1.1. Registering Link Relation Types

The link relations registry is located at <https://www.iana.org/assignments/link-relations/>. Registration requests can be made by following the instructions located there, or by sending an e-mail to the "link-relations@ietf.org" mailing list.

Registration requests consist of at least the following information:

- o ***Relation Name***: The name of the relation type
- o ***Description***: A short English description of the type's semantics. It SHOULD be stated in terms of the relationship between the link context and link target.
- o ***Reference***: Reference to the document that specifies the link relation type, preferably including a URI that can be used to retrieve a copy of the document. An indication of the relevant section(s) can also be included, but is not required.

The expert(s) can define additional fields to be collected in the registry.

General requirements for registered relation types are described in Section 2.1.1.

Registrations MUST reference a freely available, stable specification.

Note that relation types can be registered by third parties (including the expert(s)), if the expert(s) determine that an unregistered relation type is widely deployed and not likely to be registered in a timely manner otherwise. Such registrations still are subject to the requirements defined, including the need to reference a specification.

2.1.1.2. Registration Request Processing

Relation types are registered using the Specification Required policy (see Section 4.6 of [RFC8126]), which implies review and approval by a designated expert.

The goal of the registry is to reflect common use of links on the Internet. Therefore, the expert(s) should be strongly biased towards approving registrations, unless they are abusive, frivolous, not likely to be used on the Internet, or actively harmful to the Internet and/or the Web (not merely aesthetically displeasing, or architecturally dubious). As stated in Section 2.1.1, the expert(s) can withhold registration of names that are too general for the proposed application.

The expert(s) will clearly identify any issues which cause a registration to be refused. Advice about the semantics of a proposed link relation type can be given, but if it does not block registration, this should be explicitly stated.

When a request is approved, the expert(s) will inform IANA, and the registration will be processed. The IESG is the final arbiter of any objection.

2.1.2. Extension Relation Types

Applications that don't wish to register a relation type can use an extension relation type, which is a URI [RFC3986] that uniquely identifies the relation type. Although the URI can point to a resource that contains a definition of the semantics of the relation type, clients SHOULD NOT automatically access that resource to avoid overburdening its server.

The URI used for an extension relation type SHOULD be under the control of the person or party defining it, or be delegated to them.

When extension relation types are compared, they MUST be compared as strings (after converting to URIs if serialised in a different

format) in a case-insensitive fashion, character-by-character. Because of this, all-lowercase URIs SHOULD be used for extension relations.

Note that while extension relation types are required to be URIs, a serialisation of links can specify that they are expressed in another form, as long as they can be converted to URIs.

2.2. Target Attributes

Target attributes are a list of key/value pairs that describe the link or its target; for example, a media type hint.

They can be defined both by individual link relation types and by link serialisations.

This specification does not attempt to coordinate the name of target attributes, their cardinality or use. Those creating and maintaining serialisations SHOULD coordinate their target attributes to avoid conflicts in semantics or syntax, and MAY define their own registries of target attributes.

The names of target attributes SHOULD conform to the token rule, but SHOULD NOT include any of the characters "%", "'" or "*", for portability across serializations, and MUST be compared in a case-insensitive fashion.

Target attribute definitions SHOULD specify:

- o The serialisation of their values into Unicode or a subset thereof, to maximise their chances of portability across link serialisations.
- o The semantics and error handling of multiple occurrences of the target attribute on a given link.

This specification does define target attributes for use in the Link HTTP header field in Section 3.4.

3. Link Serialisation in HTTP Headers

The Link header field provides a means for serialising one or more links into HTTP headers.

The ABNF for the field value is:

```
Link           = #link-value
link-value    = "<" URI-Reference ">" *( OWS ";" OWS link-param )
link-param    = token BWS [ "=" BWS ( token / quoted-string ) ]
```

Note that any "link-param" can be generated with values using either the "token" or the "quoted-string" syntax, and therefore recipients MUST be able to parse both forms. In other words, the following parameters are equivalent:

```
x=y
x="y"
```

Previous definitions of the Link header did not equate the token and quoted-string forms explicitly; the "title" parameter was always quoted, and the "hreflang" parameter was always a token. Senders wishing to maximize interoperability will send them in those forms.

Individual "link-param"s specify their syntax in terms of the value after any necessary unquoting (as per [RFC7230], Section 3.2.6).

This specification establishes the link-params "rel", "anchor", and "rev" (which are part of the general link model), as well as "hreflang", "media", "title", "title*", and "type" (which are target attributes defined by the serialization).

3.1. Link Target

Each link-value conveys one target IRI as a URI-Reference (after conversion to one, if necessary; see [RFC3987], Section 3.1) inside angle brackets ("<>"). If the URI-Reference is relative, parsers MUST resolve it as per [RFC3986], Section 5. Note that any base IRI appearing in the message's content is not applied.

3.2. Link Context

By default, the context of a link conveyed in the Link header field is the URL of the representation it is associated with, as defined in [RFC7231], Section 3.1.4.1, serialised as a URI.

When present, the anchor parameter overrides this with another URI, such as a fragment of this resource, or a third resource (i.e., when the anchor value is an absolute URI). If the anchor parameter's value is a relative URI, parsers MUST resolve it as per [RFC3986], Section 5. Note that any base URI from the body's content is not applied.

The ABNF for the "anchor" parameter's value is:

URI-Reference ; Section 4.1 of {{RFC3986}}

Link application can choose to ignore links with an anchor parameter. For example, the application in use might not allow the link context to be assigned to a different resource. In such cases, the entire link is to be ignored; link applications MUST NOT process the link without applying the anchor.

Note that depending on HTTP status code and response headers, the link context might be "anonymous" (i.e., no link context is available). For example, this is the case on a 404 response to a GET request.

3.3. Relation Type

The relation type of a link conveyed in the Link header field is conveyed in the "rel" parameter's value. The "rel" parameter MUST be present but MUST NOT appear more than once in a given link-value; occurrences after the first MUST be ignored by parsers.

The "rel" parameter can, however, contain multiple link relation types. When this occurs, it establishes multiple links that share the same context, target, and target attributes.

The "rev" parameter has been used in the past to indicate that the semantics of the relationship are in the reverse direction. That is, a link from A to B with REL="X" expresses the same relationship as a link from B to A with REV="X". "rev" is deprecated by this specification because it often confuses authors and readers; in most cases, using a separate relation type is preferable.

The ABNF for the "rel" and "rev" parameters' values is:

```
relation-type *( 1*SP relation-type )
```

where:

```
relation-type = reg-rel-type / ext-rel-type
reg-rel-type  = LOALPHA *( LOALPHA / DIGIT / "." / "-" )
ext-rel-type  = URI ; Section 3 of {{RFC3986}}
```

Note that extension relation types are REQUIRED to be absolute URIs in Link header fields, and MUST be quoted when they contain characters not allowed in tokens, such as semicolon (";") or comma (",") (as these characters are used as delimiters in the header field itself).

3.4. Target Attributes

The Link header field defines several target attributes specific to this serialisation, and also allows extension target attributes. Target attributes are serialised in the Link header field as parameters (see [RFC7231], Section 3.1.1.1 for the definition of their syntax).

3.4.1. Serialisation-Defined Attributes

The "hreflang", "media", "title", "title*", and "type" link-params can be translated to serialisation-defined target attributes for the link.

The "hreflang" attribute, when present, is a hint indicating what the language of the result of dereferencing the link should be. Note that this is only a hint; for example, it does not override the Content-Language header field of a HTTP response obtained by actually following the link. Multiple "hreflang" attributes on a single link-value indicate that multiple languages are available from the indicated resource.

The ABNF for the "hreflang" parameter's value is:

```
Language-Tag
```

The "media" attribute, when present, is used to indicate intended destination medium or media for style information (see [W3C.REC-html5-20141028], Section 4.2.4). Its value MUST be quoted if it contains a semicolon (";") or comma (","). There MUST NOT be more than one "media" attribute in a link-value; occurrences after the first MUST be ignored by parsers.

The ABNF for the "media" parameter's value is:

```
media-query-list
```

The "title" attribute, when present, is used to label the destination of a link such that it can be used as a human-readable identifier (e.g., a menu entry) in the language indicated by the Content-Language header field (if present). The "title" attribute MUST NOT appear more than once in a given link; occurrences after the first MUST be ignored by parsers.

The "title*" link-param can be used to encode this attribute in a different character set, and/or contain language information as per [I-D.ietf-httpbis-rfc5987bis]. The "title*" link-param MUST NOT appear more than once in a given link-value; occurrences after the

first MUST be ignored by parsers. If the attribute does not contain language information, its language is indicated by the Content-Language header field (when present).

If both the "title" and "title*" link-param appear in a link, applications SHOULD use the "title*" link-param's value for the "title" attribute.

The "type" attribute, when present, is a hint indicating what the media type of the result of dereferencing the link should be. Note that this is only a hint; for example, it does not override the Content-Type header field of a HTTP response obtained by actually following the link. The "type" attribute MUST NOT appear more than once in a given link-value; occurrences after the first MUST be ignored by parsers.

The ABNF for the "type" parameter's value is:

```
type-name "/" subtype-name ; see {{RFC6838}}, Section 4.2
```

3.4.2. Extension Attributes

Other link-params are link-extensions, and are to be considered as target attributes.

Such target attributes MAY be defined to use the encoding in [I-D.ietf-httpbis-rfc5987bis] (e.g., "example" and "example*"). When both forms are present, they SHOULD be considered to be the same target attribute; applications SHOULD use the value of the name ending in "*" (after [I-D.ietf-httpbis-rfc5987bis] decoding), but MAY fall back to the other value if there is an error in decoding it, or if they do not support decoding.

3.5. Link Header Field Examples

For example:

```
Link: <http://example.com/TheBook/chapter2>; rel="previous";  
      title="previous chapter"
```

indicates that "chapter2" is previous to this resource in a logical navigation path.

Similarly,

```
Link: </>; rel="http://example.net/foo"
```

indicates that the root resource ("/") is related to this resource with the extension relation type "http://example.net/foo".

This link:

```
Link: </terms>; rel="copyright"; anchor="#foo"
```

indicates that the linked copyright terms only apply to the portion of the document indicated by the (media type-specific) fragment identifier "foo".

The example below shows an instance of the Link header field encoding multiple links, and also the use of RFC 5987 encoding to encode both non-ASCII characters and language information.

```
Link: </TheBook/chapter2>;  
      rel="previous"; title*=UTF-8'de'letztes%20Kapitel,  
      </TheBook/chapter4>;  
      rel="next"; title*=UTF-8'de'n%c3%a4chstes%20Kapitel
```

Here, both links have titles encoded in UTF-8, use the German language ("de"), and the second link contains the Unicode code point U+00E4 ("LATIN SMALL LETTER A WITH DIAERESIS").

Note that link-values can convey multiple links between the same link target and link context; for example:

```
Link: <http://example.org/>;  
      rel="start http://example.net/relation/other"
```

Here, the link to "http://example.org/" has the registered relation type "start" and the extension relation type "http://example.net/relation/other".

Finally, this header field:

```
Link: <https://example.org/>; rel="start",  
      <https://example.org/index>; rel="index"
```

is equivalent to these:

```
Link: <https://example.org/>; rel="start"  
Link: <https://example.org/index>; rel="index"
```


4. IANA Considerations

4.1. Link HTTP Header Field Registration

This specification updates the Message Header registry entry for "Link" in HTTP [RFC3864] to refer to this document.

Header field: Link
Applicable protocol: http
Status: standard
Author/change controller:
 IETF (iesg@ietf.org)
 Internet Engineering Task Force
Specification document(s):
 [this document]

4.2. Link Relation Type Registry

This specification updates the registration procedures for the Link Relation Type registry; see Section 2.1.1.1.

IANA will direct any incoming requests regarding the registry to this document and, if defined, the processes established by the expert(s); typically, this will mean referring them to the registry Web page.

Note that the expert(s) are allowed (as per Section 2.1.1.1) to define additional fields to be collected in the registry.

4.3. Link Relation Application Data Registry

This specification removes the Link Relation Application Data Registry, as it has not been used, and future use is not anticipated. IANA is instructed to remove it.

5. Security Considerations

The content of the Link header field is not secure, private or integrity-guaranteed. Use of Transport Layer Security (TLS) with HTTP ([RFC2818]) is currently the only end-to-end way to provide these properties.

Link applications ought to consider the attack vectors opened by automatically following, trusting, or otherwise using links gathered from HTTP header fields.

For example, Link header fields that use the "anchor" parameter to associate a link's context with another resource cannot be trusted since they are effectively assertions by a third party that could be

incorrect or malicious. Applications can mitigate this risk by specifying that such links should be discarded unless some relationship between the resources is established (e.g., they share the same authority).

Dereferencing links has a number of risks, depending on the application in use. For example, the Referer header [RFC7231] can expose information about the application's state (including private information) in its value. Likewise, cookies [RFC6265] are another mechanism that, if used, can become an attack vector. Applications can mitigate these risks by carefully specifying how such mechanisms should operate.

The Link header field makes extensive use of IRIs and URIs. See [RFC3987] Section 8 for security considerations relating to IRIs. See [RFC3986] Section 7 for security considerations relating to URIs. See [RFC7230] Section 9 for security considerations relating to HTTP header fields.

6. Internationalisation Considerations

Link targets may need to be converted to URIs in order to express them in serialisations that do not support IRIs. This includes the Link HTTP header field.

Similarly, the anchor parameter of the Link header field does not support IRIs, and therefore IRIs must be converted to URIs before inclusion there.

Relation types are defined as URIs, not IRIs, to aid in their comparison. It is not expected that they will be displayed to end users.

Note that registered Relation Names are required to be lower-case ASCII letters.

7. References

7.1. Normative References

[I-D.ietf-httpbis-rfc5987bis]

Reschke, J., "Indicating Character Encoding and Language for HTTP Header Field Parameters", draft-ietf-httpbis-rfc5987bis-05 (work in progress), February 2017.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3864] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", BCP 90, RFC 3864, DOI 10.17487/RFC3864, September 2004, <<http://www.rfc-editor.org/info/rfc3864>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<http://www.rfc-editor.org/info/rfc3986>>.
- [RFC3987] Duerst, M. and M. Suignard, "Internationalized Resource Identifiers (IRIs)", RFC 3987, DOI 10.17487/RFC3987, January 2005, <<http://www.rfc-editor.org/info/rfc3987>>.
- [RFC5646] Phillips, A., Ed. and M. Davis, Ed., "Tags for Identifying Languages", BCP 47, RFC 5646, DOI 10.17487/RFC5646, September 2009, <<http://www.rfc-editor.org/info/rfc5646>>.
- [RFC6838] Freed, N., Klensin, J., and T. Hansen, "Media Type Specifications and Registration Procedures", BCP 13, RFC 6838, DOI 10.17487/RFC6838, January 2013, <<http://www.rfc-editor.org/info/rfc6838>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<http://www.rfc-editor.org/info/rfc7230>>.
- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, DOI 10.17487/RFC7231, June 2014, <<http://www.rfc-editor.org/info/rfc7231>>.
- [RFC8126] 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, <<http://www.rfc-editor.org/info/rfc8126>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<http://www.rfc-editor.org/info/rfc8174>>.

[W3C.REC-css3-mediaqueries-20120619]
Rivoal, F., "Media Queries", World Wide Web Consortium
Recommendation REC-css3-mediaqueries-20120619, June 2012,
<[http://www.w3.org/TR/2012/
REC-css3-mediaqueries-20120619](http://www.w3.org/TR/2012/REC-css3-mediaqueries-20120619)>.

7.2. Informative References

[RFC2046] Freed, N. and N. Borenstein, "Multipurpose Internet Mail
Extensions (MIME) Part Two: Media Types", RFC 2046,
DOI 10.17487/RFC2046, November 1996,
<<http://www.rfc-editor.org/info/rfc2046>>.

[RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818,
DOI 10.17487/RFC2818, May 2000,
<<http://www.rfc-editor.org/info/rfc2818>>.

[RFC4287] Nottingham, M., Ed. and R. Sayre, Ed., "The Atom
Syndication Format", RFC 4287, DOI 10.17487/RFC4287,
December 2005, <<http://www.rfc-editor.org/info/rfc4287>>.

[RFC6265] Barth, A., "HTTP State Management Mechanism", RFC 6265,
DOI 10.17487/RFC6265, April 2011,
<<http://www.rfc-editor.org/info/rfc6265>>.

[W3C.REC-html5-20141028]
Hickson, I., Berjon, R., Faulkner, S., Leithead, T.,
Navara, E., O'Connor, T., and S. Pfeiffer, "HTML5",
World Wide Web Consortium Recommendation REC-
html5-20141028, October 2014,
<<http://www.w3.org/TR/2014/REC-html5-20141028>>.

Appendix A. Notes on Other Link Serialisations

Header fields (Section 3) are only one serialisation of links; other specifications have defined alternative serialisations.

A.1. Link Serialisation in HTML

HTML motivated the original syntax of the Link header field, and many of the design decisions in this document are driven by a desire to stay compatible with it.

In HTML, the link element can be mapped to links as specified here by using the "href" attribute for the target URI, and "rel" to convey the relation type, as in the Link header field. The context of the link is the URI associated with the entire HTML document. HTML also defines several attributes on links that can be seen as target attributes, including "media", "hreflang", "type" and "sizes".

Section 4.8 of HTML5 ([W3C.REC-html5-20141028]) defines modern HTML links. That document links to the Microformats Wiki as a registry; over time, the IANA registry ought to mirror its contents, and ideally eventually replace it (although that depends on the HTML community).

Surveys of existing HTML content have shown that unregistered link relation types that are not URIs are (perhaps inevitably) common. Consuming HTML implementations ought not consider such unregistered short links to be errors, but rather relation types with a local scope (i.e., their meaning is specific and perhaps private to that document).

Finally, the HTML specification gives a special meaning when the "alternate" relation type coincides with other relation types in the same link. Such links ought to be serialised in the Link header field using a single list of relation-types (e.g., rel="alternate stylesheet") to preserve this relationship.

A.2. Link Serialisation in Atom

Atom [RFC4287] is a link serialisation that conveys links in the atom:link element, with the "href" attribute indicating the link target and the "rel" attribute containing the relation type. The context of the link is either a feed locator or an entry ID, depending on where it appears; generally, feed-level links are obvious candidates for transmission as a Link header field.

When serialising an atom:link into a Link header field, it is necessary to convert link targets (if used) to URIs.

Atom defines extension relation types in terms of IRIs. This specification re-defines them as URIs, to simplify and reduce errors in their comparison.

Atom allows registered link relation types to be serialised as absolute URIs using a prefix, "http://www.iana.org/assignments/relation/". This prefix is specific to the Atom serialisation.

Furthermore, link relation types are always compared in a case-sensitive fashion; therefore, registered link relation types SHOULD be converted to their registered form (usually, lowercase) when serialised in an Atom document.

Note also that while the Link header field allows multiple relations to be serialised in a single link, atom:link does not. In this case, a single link-value may map to several atom:link elements.

As with HTML, atom:link defines some attributes that are not explicitly mirrored in the Link header field syntax, but they can also be used as link-extensions to maintain fidelity.

Appendix B. Algorithms for Parsing Link Header Fields

This appendix outlines a set of non-normative algorithms: for parsing the Link header(s) out of a header set, parsing a link header field value, and algorithms for parsing generic parts of the field value.

These algorithms are more permissive than the ABNF defining the syntax might suggest; the error handling embodied in them is a reasonable approach, but not one that is required. As such they are advisory only, and in cases where there is disagreement, the correct behaviour is defined by the body of this specification.

B.1. Parsing a Header Set for Links

This algorithm can be used to parse the Link header fields that a HTTP header set contains. Given a "header_set" of (string "field_name", string "field_value") pairs, assuming ASCII encoding, it returns a list of link objects.

1. Let "field_values" be a list containing the members of "header_set" whose "field_name" is a case-insensitive match for "link".
2. Let "links" be an empty list.
3. For each "field_value" in "field_values":

1. Let "value_links" be the result of `_Parsing A Link Field Value_` (Appendix B.2) from "field_value".
2. Append each member of "value_links" to "links".
4. Return "links".

B.2. Parsing a Link Field Value

This algorithm parses zero or more comma-separated link-values from a Link header field. Given a string "field_value", assuming ASCII encoding, it returns a list of link objects.

1. Let "links" be an empty list.
2. While "field_value" has content:
 1. Consume any leading OWS.
 2. If the first character is not "<", return "links".
 3. Discard the first character ("<").
 4. Consume up to but not including the first ">" character or end of "field_value" and let the result be "target_string".
 5. If the next character is not ">", return "links".
 6. Discard the leading ">" character.
 7. Let "link_parameters", be the result of `_Parsing Parameters_` (Appendix B.3) from "field_value" (consuming zero or more characters of it).
 8. Let "target" be the result of relatively resolving (as per [RFC3986], Section 5.2) "target_string". Note that any base URI carried in the payload body is NOT used.
 9. Let "relations_string" be the second item of the first tuple of "link_parameters" whose first item matches the string "rel", or the empty string ("") if it is not present.
 10. Split "relations_string" on RWS (removing it in the process) into a list of strings "relation_types".
 11. Let "context_string" be the second item of the first tuple of "link_parameters" whose first item matches the string "anchor". If it is not present, "context_string" is the URL

of the representation carrying the Link header [RFC7231], Section 3.1.4.1, serialised as a URI. Where the URL is anonymous, "context_string" is null.

12. Let "context" be the result of relatively resolving (as per [RFC3986], Section 5.2) "context_string", unless "context_string" is null in which case "context" is null. Note that any base URI carried in the payload body is NOT used.
13. Let "target_attributes" be an empty list.
14. For each tuple ("param_name", "param_value") of "link_parameters":
 1. If "param_name" matches "rel" or "anchor", skip this tuple.
 2. If "param_name" matches "media", "title", "title*" or "type" and "target_attributes" already contains a tuple whose first element matches the value of "param_name", skip this tuple.
 3. Append ("param_name", "param_value") to "target_attributes".
15. Let "star_param_names" be the set of "param_name"s in the ("param_name", "param_value") tuples of "link_parameters" where the last character of "param_name" is an asterisk ("*").
16. For each "star_param_name" in "star_param_names":
 1. Let "base_param_name" be "star_param_name" with the last character removed.
 2. If the implementation does not choose to support an internationalised form of a parameter named "base_param_name" for any reason (including, but not limited to, it being prohibited by the parameter's specification), remove all tuples from "link_parameters" whose first member is "star_param_name" and skip to the next "star_param_name".
 3. Remove all tuples from "link_parameters" whose first member is "base_param_name".

4. Change the first member of all tuples in "link_parameters" whose first member is "star_param_name" to "base_param_name".
17. For each "relation_type" in "relation_types":
 1. Case-normalise "relation_type" to lowercase.
 2. Append a link object to "links" with the target "target", relation type of "relation_type", context of "context", and target attributes "target_attributes".
3. Return "links".

B.3. Parsing Parameters

This algorithm parses the parameters from a header field value. Given an ASCII string "input", it returns a list of (string "parameter_name", string "parameter_value") tuples that it contains. "input" is modified to remove the parsed parameters.

1. Let "parameters" be an empty list.
2. While "input" has content:
 1. Consume any leading OWS.
 2. If the first character is not ";", return "parameters".
 3. Discard the leading ";" character.
 4. Consume any leading OWS.
 5. Consume up to but not including the first BWS, "=", ";", ",", " character or end of "input" and let the result be "parameter_name".
 6. Consume any leading BWS.
 7. If the next character is "=":
 1. Discard the leading "=" character.
 2. Consume any leading BWS.
 3. If the next character is DQUOTE, let "parameter_value" be the result of `_Parsing a Quoted String_`

(Appendix B.4) from "input" (consuming zero or more characters of it).

4. Else, consume the contents up to but not including the first ";", ",", " character or end of "input" and let the results be "parameter_value".
5. If the last character of "parameter_name" is an asterisk ("*"), decode "parameter_value" according to [I-D.ietf-httpbis-rfc5987bis]. Continue processing "input" if an unrecoverable error is encountered.
8. Else:
 1. Let "parameter_value" be an empty string.
 9. Case-normalise "parameter_name" to lowercase.
 10. Append ("parameter_name", "parameter_value") to "parameters".
 11. Consume any leading OWS.
 12. If the next character is ",", or the end of "input", stop processing "input" and return "parameters".

B.4. Parsing a Quoted String

This algorithm parses a quoted string, as per [RFC7230], Section 3.2.6. Given an ASCII string "input", it returns an unquoted string. "input" is modified to remove the parsed string.

1. Let "output" be an empty string.
2. If the first character of "input" is not DQUOTE, return "output".
3. Discard the first character.
4. While "input" has content:
 1. If the first character is a backslash ("\"):
 1. Discard the first character.
 2. If there is no more "input", return "output".
 3. Else, consume the first character and append it to "output".

2. Else, if the first character is DQUOTE, discard it and return "output".
3. Else, consume the first character and append it to "output".
5. Return "output".

Appendix C. Changes from RFC5988

This specification has the following differences from its predecessor, RFC5988:

- o The initial relation type registrations were removed, since they've already been registered by 5988.
- o The introduction has been shortened.
- o The Link Relation Application Data Registry has been removed.
- o Incorporated errata.
- o Updated references.
- o Link cardinality was clarified.
- o Terminology was changed from "target IRI" and "context IRI" to "link target" and "link context" respectively.
- o Made assigning a URI to registered relation types serialisation-specific.
- o Removed misleading statement that the link header field is semantically equivalent to HTML and Atom links.
- o More carefully defined and used "link serialisations" and "link applications."
- o Clarified the cardinality of target attributes (generically and for "type").
- o Corrected the default link context for the Link header field, to be dependent upon the identity of the representation (as per RFC7231).
- o Defined a suggested parsing algorithm for the Link header.
- o The value space of target attributes and their definition has been specified.

- o The ABNF has been updated to be compatible with [RFC7230]. In particular, whitespace is now explicit.
- o Some parameters on the HTTP header field can now appear as a token.
- o Parameters on the HTTP header can now be value-less.
- o Handling of quoted strings is now defined by [RFC7230].
- o The "type" header field parameter now needs to be quoted (as "token" does not allow "/").

Author's Address

Mark Nottingham

E-Mail: mnot@mnot.net

URI: <https://www.mnot.net/>

DISPATCH
Internet-Draft
Intended status: Standards Track
Expires: September 14, 2017

M. Petit-Huguenin
Impedance Mismatch
G. Salgueiro
Cisco
March 13, 2017

Path MTU Discovery (PMTUD) for RTP/RTCP
draft-petithuguenin-dispatch-rtp-pmtud-00

Abstract

This document describes an implementation of the Path MTU Discovery (PMTUD) protocol for RTP sessions.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 14, 2017.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Overview of Operations	2
3. Terminology	2
4. Probe Support Signaling	3
5. Path MTU Discovery Using the Simple Probing Mechanism	3
6. Path MTU Discovery Using the Complete Probing Mechanism	3
7. Security Considerations	4
8. IANA Considerations	4
9. References	4
9.1. Normative References	4
9.2. Informative References	5
Authors' Addresses	5

1. Introduction

The Guidelines for Writers of RTP Payload Formats (RFC 2736, BCP 36 [RFC2736]) states in Section 4 that "[i]f a codec's frame size is larger than the MTU, the payload format must not rely on IP fragmentation." Similarly, RFC 3550 [RFC3550] states that "...only the subset [of RR packets into one compound RTCP packet] that will fit into one MTU SHOULD be included in each interval."

These statements can be extended to the Path MTU, as fragmentation along the media path is no better than fragmentation on the first link-layer.

RTP and RTCP [RFC3550] were not designed with a mechanism to discover the Path MTU, so this document describes a way to add this capability by using the PMTUD protocol defined in [I-D.ietf-tram-stun-pmtud].

2. Overview of Operations

Multiplexing between RTP/RTCP packets and STUN packets is a well-known technique used for example to discover the IP address of a NAT [RFC5389] or to check connectivity [RFC5245].

The PMTUD mechanism for RTP/RTCP uses either the Simple Probing Mechanism described in Section 4.1 of [I-D.ietf-tram-stun-pmtud] or the Complete Probing Mechanism described in Section 4.2.

3. 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 [RFC2119].

4. Probe Support Signaling

Real-time media protocols (SIP [RFC3261], Jingle [XEP-0166]) that are using the Offer/Answer protocol [RFC3264] signals their support of this specification by the usage of an "a:x-pmtud" attribute in the SDP. This attribute can be used at the session-level or at the media-level.

An Offerer indicates the support of this specification by adding an "a:x-pmtud" attribute in the SDP sent. An Answerer receiving an SDP containing an "a:x-pmtud" attribute and supporting this specification can immediately start probing for the PMTU, as described in Section 5.2 of [I-D.ietf-tram-stun-pmtud]. Even if the SDP received by an Answerer does not contain an "a:x-pmtud" attribute, an Answerer supporting this specification MUST insert an "a:x-pmtud" attribute in the SDP it will send.

Realtime media protocols that support ICE [RFC5245] (i.e. WebRTC, in addition to the protocols listed above) MAY signal that a specific candidate support for this specification differs from what is declared at the session-level or media-level of the SDP by inserting an extension attribute with values "pmtud on" or "pmtud off" in the candidate line.

5. Path MTU Discovery Using the Simple Probing Mechanism

When initiating the Probe transactions, as described in Section 4.1 of [I-D.ietf-tram-stun-pmtud], the RTP/RTCP client MUST use the same IP address and port destination that are used as the destination for the RTP or RTCP packets.

The server side MUST be prepared to demultiplex the Probe Requests from the RTP/RTCP packets and other STUN messages.

6. Path MTU Discovery Using the Complete Probing Mechanism

When sending the Probe Indications, the RTP/RTCP client MUST use the same source IP address and port and same IP address and port destination that are used for the RTP or RTCP packets.

Any STUN message sent along the RTP/RTCP packets, like ICE connectivity checks, media keep-alive, or consent packets MUST be used to populate the identifier list described in Section 4.2.3 of [I-D.ietf-tram-stun-pmtud].

For a STUN message, the identifier is made up of the first 12 bytes of the Transaction ID.

For an RTP packet, the identifier is made up of the SSRC concatenated with the Sequence Number, for a total of 12 bytes.

For an RTCP packet, the identifier is made up of the Reporter SSRC concatenated with the last 4 bytes of the Extended Highest Sequence Number Received, for a total of 12 bytes.

7. Security Considerations

TBD.

8. IANA Considerations

TBD

9. References

9.1. Normative References

- [I-D.ietf-tram-stun-pmtud]
Petit-Huguenin, M. and G. Salgueiro, "Path MTU Discovery Using Session Traversal Utilities for NAT (STUN)", draft-ietf-tram-stun-pmtud-05 (work in progress), February 2017.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, DOI 10.17487/RFC3264, June 2002, <<http://www.rfc-editor.org/info/rfc3264>>.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, DOI 10.17487/RFC3550, July 2003, <<http://www.rfc-editor.org/info/rfc3550>>.
- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, DOI 10.17487/RFC5245, April 2010, <<http://www.rfc-editor.org/info/rfc5245>>.

[RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing,
"Session Traversal Utilities for NAT (STUN)", RFC 5389,
DOI 10.17487/RFC5389, October 2008,
<<http://www.rfc-editor.org/info/rfc5389>>.

9.2. Informative References

[RFC2736] Handley, M. and C. Perkins, "Guidelines for Writers of RTP
Payload Format Specifications", BCP 36, RFC 2736,
DOI 10.17487/RFC2736, December 1999,
<<http://www.rfc-editor.org/info/rfc2736>>.

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston,
A., Peterson, J., Sparks, R., Handley, M., and E.
Schooler, "SIP: Session Initiation Protocol", RFC 3261,
DOI 10.17487/RFC3261, June 2002,
<<http://www.rfc-editor.org/info/rfc3261>>.

Authors' Addresses

Marc Petit-Huguenin
Impedance Mismatch

Email: marc@petit-huguenin.org

Gonzalo Salgueiro
Cisco Systems, Inc.
7200-12 Kit Creek Road
Research Triangle Park, NC 27709
United States

Email: gsalguei@cisco.com