Well-Known Uniform Resource Identifiers (URIs)
draft-nottingham-rfc5785bis-11

Abstract

   This memo defines a path prefix for "well-known locations", "/.well-
   known/", in selected Uniform Resource Identifier (URI) schemes.

   In doing so, it obsoletes RFC 5785, and updates the URI schemes
   defined in RFC 7230 to reserve that space. It also updates RFC 7595
   to track URI schemes that support well-known URIs in their registry.

Note to Readers

   RFC EDITOR: please remove this section before publication

   This draft is a proposed revision of RFC5875.

   The issues list for this draft can be found at
   https://github.com/mnot/I-D/labels/rfc5785bis [1].

   The most recent (often, unpublished) draft is at
   https://mnot.github.io/I-D/rfc5785bis/ [2].

   Recent changes are listed at https://github.com/mnot/I-D/commits/gh-
   pages/rfc5785bis [3].

   See also the draft's current status in the IETF datatracker, at
   https://datatracker.ietf.org/doc/draft-nottingham-rfc5785bis/ [4].

Status of This Memo

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 October 10, 2019.

Copyright Notice

Table of Contents

1. Introduction

Some applications on the Web require the discovery of information
about an origin [RFC6454] (sometimes called "site-wide metadata")
before making a request. For example, the Robots Exclusion Protocol
(http://www.robotstxt.org/ [5]) specifies a way for automated
processes to obtain permission to access resources; likewise, the
Platform for Privacy Preferences [P3P] tells user-agents how to
discover privacy policy before interacting with an origin server.

While there are several ways to access per-resource metadata (e.g.,
HTTP header fields, WebDAV's PROPFIND [RFC4918]), the perceived
overhead (either in terms of client-perceived latency and/or
deployment difficulties) associated with them often precludes their
use in these scenarios.

At the same time, it has become more popular to use HTTP as a
substrate for non-Web protocols. Sometimes, such protocols need a
way to locate one or more resources on a given host.

When this happens, one solution is to designate a "well-known
location" for data or services related to the origin overall, so that
it can be easily located. However, this approach has the drawback of
risking collisions, both with other such designated "well-known
locations" and with resources that the origin has created (or wishes
to create). Furthermore, defining well-known locations usurp's the
origin's control over its own URI space [RFC7320].

To address these uses, this memo reserves a path prefix in HTTP,
HTTPS, WS and WSS URIs for these "well-known locations", "/.well-
known/". Future specifications that need to define a resource for
such metadata can register their use to avoid collisions and minimise
impingement upon origins' URI space.

Well-known URIs can also be used with other URI schemes, but only
when those schemes' definitions explicitly allow it.

2. 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.

3. **Well-Known URIs**

   A well-known URI is a URI [RFC3986] whose path component begins with
   the characters "/.well-known/", provided that the scheme is
   explicitly defined to support well-known URIs.

   For example, if an application registers the name 'example', the
   corresponding well-known URI on 'http://www.example.com/' would be
   'http://www.example.com/.well-known/example'.

   This specification updates the "http" [RFC7230] and "https" [RFC7230]
   schemes to support well-known URIs. Other existing schemes can use
   the appropriate process for updating their definitions; for example,
   [RFC8307] does so for the "ws" and "wss" schemes. The Uniform
   Resource Identifier (URI) Schemes Registry tracks which schemes
   support well-known URIs; see Section 5.2.

   Applications that wish to mint new well-known URIs MUST register
   them, following the procedures in Section 5.1, subject to the
   following requirements.

   Registered names MUST conform to the segment-nz production in
   [RFC3986]. This means they cannot contain the "/" character.

   Registered names for a specific application SHOULD be correspondingly
   precise; "squatting" on generic terms is not encouraged. For
   example, if the Example application wants a well-known location for
   metadata, an appropriate registered name might be "example-metadata"
   or even "example.com-metadata", not "metadata".

   At a minimum, a registration will reference a specification that
   defines the format and associated media type(s) to be obtained by
   dereferencing the well-known URI, along with the URI scheme(s) that
   the well-known URI can be used with. If no URI schemes are
   explicitly specified, "http" and "https" are assumed.

   Typically, applications will use the default port for the given
   scheme; if an alternative port is used, it MUST be explicitly
   specified by the application in question.

   Registrations MAY also contain additional information, such as the
   syntax of additional path components, query strings and/or fragment
   identifiers to be appended to the well-known URI, or protocol-
   specific details (e.g., HTTP [RFC7231] method handling).

   Note that this specification defines neither how to determine the
   hostname to use to find the well-known URI for a particular
   application, nor the scope of the metadata discovered by

dereferencing the well-known URI; both should be defined by the
application itself.

Also, this specification does not define a format or media-type for
the resource located at "/.well-known/" and clients should not expect
a resource to exist at that location.

Well-known URIs are rooted in the top of the path's hierarchy; they
are not well-known by definition in other parts of the path. For
example, "/.well-known/example" is a well-known URI, whereas
"/foo/.well-known/example" is not.

See also Section 4 for Security Considerations regarding well-known
locations.

3.1. Registering Well-Known URIs

The "Well-Known URIs" registry is located at
"https://www.iana.org/assignments/well-known-uris/". Registration
requests can be made by following the instructions located there or
by sending an email to the "wellknown-uri-review@ietf.org" mailing
list.

Registration requests consist of at least the following information:

URI suffix: The name requested for the well-known URI, relative to
   "/.well-known/"; e.g., "example".

Change controller: For Standards-Track RFCs, state "IETF". For
   others, give the name of the responsible party. Other details
   (e.g., e-mail address, home page URI) may also be included.

Specification document(s): Reference to the document that specifies
   the field, preferably including a URI that can be used to retrieve
   a copy of the document. An indication of the relevant sections
   may also be included, but is not required.

Status: One of "permanent" or "provisional". See guidance below.

Related information: Optionally, citations to additional documents
   containing further relevant information.

General requirements for registered values are described in
Section 3.

Values defined by standards-track RFCs and other open standards (in
the sense of [RFC2026], Section 7.1.1) have a status of "permanent".
Other values can also be registered as permanent, if the Experts find

that they are in use, in consultation with the community. Other
values should be registered as "provisional".

Provisional entries can be removed by the Experts if - in
consultation with the community - the Experts find that they are not
in use. The Experts can change a provisional entry's status to
permanent; in doing so, the Experts should consider how widely used a
value is, and consult the community beforehand.

Note that "consult with the community" above refers to those
responsible for the URI scheme(s) in question. Generally, this would
take place on the mailing list(s) of the appropriate Working Group(s)
(possibly historical), or on art@ietf.org if no such list exists.

Well-known URIs can be registered by third parties (including the
expert(s)), if the expert(s) determine that an unregistered well-
known URI 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.

[4](). Security Considerations

Applications minting new well-known URIs, as well as administrators
deploying them, will need to consider several security-related
issues, including (but not limited to) exposure of sensitive data,
denial-of-service attacks (in addition to normal load issues), server
and client authentication, vulnerability to DNS rebinding attacks,
and attacks where limited access to a server grants the ability to
affect how well-known URIs are served.

[RFC3552] contains some examples of potential security considerations
that may be relevant to application protocols and administrators
deploying them.

[4.1](). Protecting Well-Known Resources

Because well-known locations effectively represent the entire origin,
server operators should appropriately control the ability to write to
them. This is especially true when more than one entity is co-
located on the same origin. Even for origins that are controlled by
and represent a single entity, due care should be taken to assure
that write access to well-known locations is not granted unwittingly,
either externally through server configuration, or locally through
implementation permissions (e.g., on a filesystem).

4.2. Interaction with Web Browsing

   Applications using well-known URIs for "http" or "https" URLs need to
   be aware that well-known resources will be accessible to Web
   browsers, and therefore are able to be manipulated by content
   obtained from other parts of that origin. If an attacker is able to
   inject content (e.g., through a Cross-Site Scripting vulnerability),
   they will be able to make potentially arbitrary requests to the well-
   known resource.

   HTTP and HTTPS also use origins as a security boundary for many other
   mechanisms, including (but not limited to) Cookies [RFC6265], Web
   Storage [WEBSTORAGE] and many capabilities.

   Applications defining well-known locations should not assume that
   they have sole access to these mechanisms, or that they are the only
   application using the origin. Depending on the nature of the
   application, mitigations can include:

   o Encrypting sensitive information

   o Allowing flexibility in the use of identifiers (e.g., Cookie
      names) to avoid collisions with other applications

   o Using the 'HttpOnly' flag on Cookies to assure that cookies are
      not exposed to browser scripting languages [RFC6265]

   o Using the 'Path' parameter on Cookies to assure that they are not
      available to other parts of the origin [RFC6265]

   o Using X-Content-Type-Options: nosniff [FETCH] to assure that
      content under attacker control can't be coaxed into a form that is
      interpreted as active content by a Web browser

   Other good practices include:

   o Using an application-specific media type in the Content-Type
      header field, and requiring clients to fail if it is not used

   o Using Content-Security-Policy [CSP] to constrain the capabilities
      of active content (such as HTML [HTML5]), thereby mitigating
      Cross-Site Scripting attacks

   o Using Referrer-Policy [REFERRER-POLICY] to prevent sensitive data
      in URLs from being leaked in the Referer request header field

   o Avoiding use of compression on any sensitive information (e.g.,
      authentication tokens, passwords), as the scripting environment

offered by Web browsers allows an attacker to repeatedly probe the compression space; if the attacker has access to the path of the communication, they can use this capability to recover that information.

4.3. Scoping Applications

This memo does not specify the scope of applicability for the information obtained from a well-known URI, and does not specify how to discover a well-known URI for a particular application.

Individual applications using this mechanism must define both aspects; if this is not specified, security issues can arise from implementation deviations and confusion about boundaries between applications.

Applying metadata discovered in a well-known URI to resources other than those co-located on the same origin risks administrative as well as security issues. For example, allowing "https://example.com/.well-known/example" to apply policy to "https://department.example.com", "https://www.example.com" or even "https://www.example.com:8000" assumes a relationship between hosts where there might be none, giving control to a potential attacker.

Likewise, specifying that a well-known URI on a particular hostname is to be used to bootstrap a protocol can cause a large number of undesired requests. For example, if a well-known HTTPS URI is used to find policy about a separate service such as e-mail, it can result in a flood of requests to Web servers, even if they don't implement the well-known URI. Such undesired requests can resemble a denial-of-services attack.

4.4. Hidden Capabilities

Applications using well-known locations should consider that some server administrators might be unaware of its existence (especially on operating systems that hide directories whose names begin with "."). This means that if an attacker has write access to the .well-known directory, they would be able to control its contents, possibly without the administrator realising it.

5. IANA Considerations

5.1. The Well-Known URI Registry

This specification updates the registration procedures for the "Well-Known URI" registry, first defined in [RFC5785]; see Section 3.1.

Well-known URIs are registered on the advice of one or more Experts, with a Specification Required (using terminology from [RFC8126]).

The Experts' primary considerations in evaluating registration requests are:

o Conformance to the requirements in Section 3

o The availability and stability of the specifying document

o The considerations outlined in Section 4

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.

Upon publication, IANA should:

o Update the status of all existing registrations to "permanent".

5.2. The Uniform Resource Identifier (URI) Schemes Registry

This specification adds a field to the registration template of the Uniform Resource Identifier (URI) Schemes Registry, with the name "Well-Known URI Support" and a default value of "-".

If a URI scheme explicitly has been specified to use well-known URIs as per Section 3, the value changes to a reference to that specification. Initial values not equal to "-" are given in Table 1.

```
      +------------+------------------------+
      | URI Scheme | Well-Known URI Support |
      +------------+------------------------+
      | coap       | [RFC7252]              |
      | coap+tcp   | [RFC8323]              |
      | coap+ws    | [RFC8323]              |
      | coaps      | [RFC7252]              |
      | coaps+tcp  | [RFC8323]              |
      | coaps+ws   | [RFC8323]              |
      | http       | [this document]        |
      | https      | [this document]        |
      | ws         | [RFC8307]              |
      | wss        | [RFC8307]              |
      +------------+------------------------+
```

Table 1: Rows in URI scheme registry with nonempty new column

6. References

6.1. 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,
              <https://www.rfc-editor.org/info/rfc2119>.

   [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform
              Resource Identifier (URI): Generic Syntax", STD 66,
              RFC 3986, DOI 10.17487/RFC3986, January 2005,
              <https://www.rfc-editor.org/info/rfc3986>.

   [RFC6454] Barth, A., "The Web Origin Concept", RFC 6454,
              DOI 10.17487/RFC6454, December 2011,
              <https://www.rfc-editor.org/info/rfc6454>.

   [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,
              <https://www.rfc-editor.org/info/rfc7230>.

   [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,
              <https://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, <https://www.rfc-editor.org/info/rfc8174>.

6.2. Informative References

   [CSP]   West, M., "Content Security Policy Level 3", World Wide
              Web Consortium WD WD-CSP3-20160913, September 2016,
              <https://www.w3.org/TR/2016/WD-CSP3-20160913>.

   [FETCH]   WHATWG, "Fetch - Living Standard", n.d.,
              <https://fetch.spec.whatwg.org>.

   [HTML5]   WHATWG, "HTML - Living Standard", n.d.,
              <https://html.spec.whatwg.org>.

   [P3P]   Marchiori, M., "The Platform for Privacy Preferences 1.0
              (P3P1.0) Specification", World Wide Web Consortium
              Recommendation REC-P3P-20020416, April 2002,
              <http://www.w3.org/TR/2002/REC-P3P-20020416>.

[REFERRER-POLICY]
          Eisinger, J. and E. Stark, "Referrer Policy", World Wide
          Web Consortium CR CR-referrer-policy-20170126, January
          2017,
          <https://www.w3.org/TR/2017/CR-referrer-policy-20170126>.

[RFC2026] Bradner, S., "The Internet Standards Process -- Revision
          3", BCP 9, RFC 2026, DOI 10.17487/RFC2026, October 1996,
          <https://www.rfc-editor.org/info/rfc2026>.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC
          Text on Security Considerations", BCP 72, RFC 3552,
          DOI 10.17487/RFC3552, July 2003,
          <https://www.rfc-editor.org/info/rfc3552>.

[RFC4918] Dusseault, L., Ed., "HTTP Extensions for Web Distributed
          Authoring and Versioning (WebDAV)", RFC 4918,
          DOI 10.17487/RFC4918, June 2007,
          <https://www.rfc-editor.org/info/rfc4918>.

[RFC5785] Nottingham, M. and E. Hammer-Lahav, "Defining Well-Known
          Uniform Resource Identifiers (URIs)", RFC 5785,
          DOI 10.17487/RFC5785, April 2010,
          <https://www.rfc-editor.org/info/rfc5785>.

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

[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,
          <https://www.rfc-editor.org/info/rfc7231>.

[RFC7320] Nottingham, M., "URI Design and Ownership", BCP 190,
          RFC 7320, DOI 10.17487/RFC7320, July 2014,
          <https://www.rfc-editor.org/info/rfc7320>.

[RFC8307] Bormann, C., "Well-Known URIs for the WebSocket Protocol",
          RFC 8307, DOI 10.17487/RFC8307, January 2018,
          <https://www.rfc-editor.org/info/rfc8307>.

[WEBSTORAGE]
          Hickson, I., "Web Storage (Second Edition)", World Wide
          Web Consortium Recommendation REC-webstorage-20160419,
          April 2016,
          <http://www.w3.org/TR/2016/REC-webstorage-20160419>.

6.3. URIs

   [1] https://github.com/mnot/I-D/labels/rfc5785bis

   [2] https://mnot.github.io/I-D/rfc5785bis/

   [3] https://github.com/mnot/I-D/commits/gh-pages/rfc5785bis

   [4] https://datatracker.ietf.org/doc/draft-nottingham-rfc5785bis/

   [5] http://www.robotstxt.org/

Appendix A. Frequently Asked Questions

   Aren't well-known locations bad for the Web? They are, but for
      various reasons - both technical and social - they are sometimes
      necessary. This memo defines a "sandbox" for them, to reduce the
      risks of collision and to minimise the impact upon pre-existing
      URIs on sites.

   Why /.well-known? It's short, descriptive, and according to search
      indices, not widely used.

   What impact does this have on existing mechanisms, such as P3P and
   robots.txt?
      None, until they choose to use this mechanism.

   Why aren't per-directory well-known locations defined? Allowing
      every URI path segment to have a well-known location (e.g.,
      "/images/.well-known/") would increase the risks of colliding with
      a pre-existing URI on a site, and generally these solutions are
      found not to scale well, because they're too "chatty".

Appendix B. Changes from RFC5785

   o Allow non-Web well-known locations

   o Adjust IANA instructions

   o Update references

   o Various other clarifications

   o Track supporting schemes in the URI Scheme registry

Author's Address

Mark Nottingham

Email: mnot@mnot.net
URI: https://www.mnot.net/