

Network Working Group  
Internet-Draft  
Obsoletes: 8615 (if approved)  
Intended status: Standards Track  
Expires: 27 August 2026

M. Nottingham  
23 February 2026

Well-Known Uniform Resource Identifiers (URIs)  
draft-nottingham-rfc8615bis-02

## 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 8615.

## Status of This Memo

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

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

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

This Internet-Draft will expire on 27 August 2026.

## Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the  
document authors. All rights reserved.

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

## Table of Contents

1. Introduction . . . . .	2
2. Notational Conventions . . . . .	3
3. Well-Known URIs . . . . .	3
3.1. Registering Well-Known URIs . . . . .	4
4. Security Considerations . . . . .	6
4.1. Protecting Well-Known Resources . . . . .	6
4.2. Interaction with Web Browsing . . . . .	6
4.3. Scoping Applications . . . . .	7
4.4. Hidden Capabilities . . . . .	8
5. IANA Considerations . . . . .	8
5.1. The Well-Known URI Registry . . . . .	8
6. References . . . . .	9
6.1. Normative References . . . . .	9
6.2. Informative References . . . . .	9
Appendix A. Frequently Asked Questions . . . . .	11
Appendix B. Changes from RFC 8615 . . . . .	11
Author's Address . . . . .	11

## 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 [RFC9309] 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, PROPFIND in Web Distributed Authoring and Versioning (WebDAV) [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 usurps the origin's control over its own URI space [RFC7320].

To address these uses, this memo reserves a path prefix in nominated URI schemes 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'.

Previous revisions of this specification updated the "http" [RFC9110] and "https" [RFC9110] 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.

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 [RFC9110] 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". Syntactic requirements are described in Section 3.

Change controller: For Standards Track RFCs, state "IETF". For

others, give the name of the responsible party. Other details (e.g., email 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. Specification documents are expected to address the security considerations in Section 4.

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

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

Values defined by Standards Track RFCs and other open standards (in the sense of [RFC2026], Section 7.1.1) can be registered with a status of "permanent".

The change controller of a permanent registration MUST be the responsible party for the specification document(s). The change controller of a permanent registration can request that its status be changed to "obsoleted" or "deprecated" as circumstances require. If a change controller becomes defunct or cannot be contacted, the experts can change a registration's status as required, in consultation with the community.

Values not defined by open standards processes can be registered with a status of "provisional".

A request for provisional registration can be denied by the experts if they find that it uses a URI suffix that is overly generic or of potential community value in the future. For example, a provisional registration request for a commonly used term can be denied on the basis that it could preempt a future standards effort's use of the term, even if no such effort is yet in progress.

Provisional entries can be promoted to "permanent" by the experts if -- in consultation with the community -- the experts find they are in broad use.

Provisional entries can be removed by the experts if -- in consultation with the community -- the experts find that they are not in broad use one year or more after registration.

Provisional entries can be removed by the experts at any time if the specification document(s) become unavailable, deprecated, obsoleted, or similar. In the case that a document becomes unavailable, the experts must attempt contact with the change controller to clarify its status first.

Note that "consult 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 concluded), or on <art@ietf.org> if no such list exists.

The experts will keep public records of their decisions. Decisions by the experts can be appealed to the IESG.

#### 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 various capabilities.

An application that defines well-known locations should not assume that it has sole access to these mechanisms or that it is the only application using the origin. Depending on the nature of the application, mitigations can include:

- \* Encrypting sensitive information
- \* Allowing flexibility in the use of identifiers (e.g., cookie names) to avoid collisions with other applications
- \* Using the 'HttpOnly' flag on cookies to assure that cookies are not exposed to browser scripting languages [RFC6265]
- \* Using the 'Path' parameter on cookies to assure that they are not available to other parts of the origin [RFC6265]
- \* Using X-Content-Type-Options: nosniff [FETCH] to assure that content under attacker control can't be coerced into a form that is interpreted as active content by a Web browser

Other good practices include:

- \* Using an application-specific media type in the Content-Type header field, and requiring clients to fail if it is not used
- \* Using Content-Security-Policy [CSP] to constrain the capabilities of active content (such as HTML [HTML5]), thereby mitigating Cross-Site Scripting attacks
- \* Using Referrer-Policy [REFERRER-POLICY] to prevent sensitive data in URLs from being leaked in the Referer request header field
- \* 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, thereby 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 email, 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-service attack.

#### 4.4. Hidden Capabilities

Applications using well-known locations should consider that some server administrators might be unaware of their 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].

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

See Section 3.1 for the registration request procedure.

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

## 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/rfc/rfc2119>>.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <<https://www.rfc-editor.org/rfc/rfc3986>>.
- [RFC6454] Barth, A., "The Web Origin Concept", RFC 6454, DOI 10.17487/RFC6454, December 2011, <<https://www.rfc-editor.org/rfc/rfc6454>>.
- [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/rfc/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/rfc/rfc8174>>.
- [RFC9110] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP Semantics", STD 97, RFC 9110, DOI 10.17487/RFC9110, June 2022, <<https://www.rfc-editor.org/rfc/rfc9110>>.

### 6.2. Informative References

- [CSP] "Content Security Policy Level 3", W3C WD WD-CSP3-20160913, W3C WD-CSP3-20160913, 13 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., Ed., "The Platform for Privacy Preferences 1.0 (P3P1.0) Specification", W3C REC REC-P3P-20020416, W3C REC-P3P-20020416, 16 April 2002, <<https://www.w3.org/TR/2002/REC-P3P-20020416/>>.

## [REFERRER-POLICY]

Stark, E., Ed. and J. Eisinger, Ed., "Referrer Policy",  
W3C CR CR-referrer-policy-20170126, W3C CR-referrer-  
policy-20170126, 26 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/rfc/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/rfc/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/rfc/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/rfc/rfc5785>>.

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

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

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

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

[RFC8323] Bormann, C., Lemay, S., Tschofenig, H., Hartke, K.,  
Silverajan, B., and B. Raymor, Ed., "CoAP (Constrained  
Application Protocol) over TCP, TLS, and WebSockets",  
RFC 8323, DOI 10.17487/RFC8323, February 2018,  
<<https://www.rfc-editor.org/rfc/rfc8323>>.

[RFC9309] Koster, M., Illyes, G., Zeller, H., and L. Sassman,  
"Robots Exclusion Protocol", RFC 9309,  
DOI 10.17487/RFC9309, September 2022,  
<<https://www.rfc-editor.org/rfc/rfc9309>>.

[WEBSTORAGE]

Hickson, I., Ed., "Web Storage (Second Edition)", W3C REC  
REC-webstorage-20160419, W3C REC-webstorage-20160419, 19  
April 2016,  
<<https://www.w3.org/TR/2016/REC-webstorage-20160419/>>.

## 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 preexisting 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 preexisting URI on a site, and generally these solutions are found not to scale well because they're too "chatty".

## Appendix B. Changes from RFC 8615

- \* Updated registration guidelines

## Author's Address

Mark Nottingham  
Email: [mnot@mnot.net](mailto:mnot@mnot.net)  
URI: <https://www.mnot.net/>