

Internet Engineering Task Force (IETF)  
Request for Comments: 5785  
Updates: 2616, 2818  
Category: Standards Track  
ISSN: 2070-1721

M. Nottingham  
E. Hammer-Lahav  
April 2010

## Defining Well-Known Uniform Resource Identifiers (URIs)

### Abstract

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

### Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5785>.

### Copyright Notice

Copyright (c) 2010 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
1.1. Appropriate Use of Well-Known URIs . . . . .	3
2. Notational Conventions . . . . .	3
3. Well-Known URIs . . . . .	3
4. Security Considerations . . . . .	4
5. IANA Considerations . . . . .	4
5.1. The Well-Known URI Registry . . . . .	4
5.1.1. Registration Template . . . . .	5
6. References . . . . .	5
6.1. Normative References . . . . .	5
6.2. Informative References . . . . .	5
Appendix A. Acknowledgements . . . . .	7
Appendix B. Frequently Asked Questions . . . . .	7

## 1. Introduction

It is increasingly common for Web-based protocols to require the discovery of policy or other information about a host ("site-wide metadata") before making a request. For example, the Robots Exclusion Protocol <<http://www.robotstxt.org/>> specifies a way for automated processes to obtain permission to access resources; likewise, the Platform for Privacy Preferences [W3C.REC-P3P-20020416] tells user-agents how to discover privacy policy beforehand.

While there are several ways to access per-resource metadata (e.g., HTTP headers, 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.

When this happens, it is common to designate a "well-known location" for such data, 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 pre-existing resources.

To address this, this memo defines a path prefix in HTTP(S) URIs for these "well-known locations", `"/.well-known/`". Future specifications that need to define a resource for such site-wide metadata can register their use to avoid collisions and minimise impingement upon sites' URI space.

### 1.1. Appropriate Use of Well-Known URIs

There are a number of possible ways that applications could use Well-known URIs. However, in keeping with the Architecture of the World-Wide Web [W3C.REC-webarch-20041215], well-known URIs are not intended for general information retrieval or establishment of large URI namespaces on the Web. Rather, they are designed to facilitate discovery of information on a site when it isn't practical to use other mechanisms; for example, when discovering policy that needs to be evaluated before a resource is accessed, or when using multiple round-trips is judged detrimental to performance.

As such, the well-known URI space was created with the expectation that it will be used to make site-wide policy information and other metadata available directly (if sufficiently concise), or provide references to other URIs that provide such metadata.

## 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 RFC 2119 [RFC2119].

## 3. Well-Known URIs

A well-known URI is a URI [RFC3986] whose path component begins with the characters `/.well-known/`, and whose scheme is `HTTP`, `HTTPS`, or another scheme that has explicitly been specified to use well-known URIs.

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

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'`.

Registered names **MUST** conform to the segment-nz production in [RFC3986].

Note that this specification defines neither how to determine the authority to use for a particular context, nor the scope of the metadata discovered by dereferencing the well-known URI; both should be defined by the application itself.

Typically, a registration will reference a specification that defines the format and associated media type to be obtained by dereferencing the well-known URI.

It 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 [RFC2616] method handling).

Note that 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.

#### 4. Security Considerations

This memo does not specify the scope of applicability of metadata or policy 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.

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.

#### 5. IANA Considerations

##### 5.1. The Well-Known URI Registry

This document establishes the well-known URI registry.

Well-known URIs are registered on the advice of one or more Designated Experts (appointed by the IESG or their delegate), with a Specification Required (using terminology from [RFC5226]). However, to allow for the allocation of values prior to publication, the Designated Expert(s) may approve registration once they are satisfied that such a specification will be published.

Registration requests should be sent to the `wellknown-uri-review@ietf.org` mailing list for review and comment, with an appropriate subject (e.g., "Request for well-known URI: example").

Before a period of 14 days has passed, the Designated Expert(s) will either approve or deny the registration request, communicating this decision both to the review list and to IANA. Denials should include an explanation and, if applicable, suggestions as to how to make the

request successful. Registration requests that are undetermined for a period longer than 21 days can be brought to the IESG's attention (using the `iesg@iesg.org` mailing list) for resolution.

#### 5.1.1. Registration Template

**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., postal address, 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.

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

## 6. References

### 6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

### 6.2. Informative References

- [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, June 1999.
- [RFC4918] Dusseault, L., "HTTP Extensions for Web Distributed Authoring and Versioning (WebDAV)", RFC 4918, June 2007.

[W3C.REC-P3P-20020416]

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](http://www.w3.org/TR/2002/REC-P3P-20020416)>.

[W3C.REC-webarch-20041215]

Jacobs, I. and N. Walsh, "Architecture of the World Wide Web, Volume One", World Wide Web Consortium Recommendation REC- webarch-20041215, December 2004, <[http:// www.w3.org/TR/2004/REC-webarch-20041215](http://www.w3.org/TR/2004/REC-webarch-20041215)>.

## Appendix A. Acknowledgements

We would like to acknowledge the contributions of everyone who provided feedback and use cases for this document; in particular, Phil Archer, Dirk Balfanz, Adam Barth, Tim Bray, Brian Eaton, Brad Fitzpatrick, Joe Gregorio, Paul Hoffman, Barry Leiba, Ashok Malhotra, Breno de Medeiros, John Panzer, and Drummond Reed. However, they are not responsible for errors and omissions.

## Appendix B. Frequently Asked Questions

### 1. Aren't well-known locations bad for the Web?

They are, but for various reasons -- both technical and social -- they are commonly used and their use is increasing. This memo defines a "sandbox" for them, to reduce the risks of collision and to minimise the impact upon pre-existing URIs on sites.

### 2. Why /.well-known?

It's short, descriptive, and according to search indices, not widely used.

### 3. What impact does this have on existing mechanisms, such as P3P and robots.txt?

None, until they choose to use this mechanism.

### 4. 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".

Authors' Addresses

Mark Nottingham

E-Mail: [mnot@mnot.net](mailto:mnot@mnot.net)

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

Eran Hammer-Lahav

E-Mail: [eran@hueniverse.com](mailto:eran@hueniverse.com)

URI: <http://hueniverse.com/>