

5785bis

IETF101, London

- RFC5785 introduced “Well-Known URIs”
- Use cases:
 - robots.txt
 - P3P
 - DNT

- URI paths are under the control of their authority (i.e., the origin's administrator). Cf: Architecture of the WWW, Vol1
- Standards should not encroach into this space.
- BCP190 explains why.
- Well-Known URIs are a very limited carve-out for standards to use for metadata **about the origin**.

Meanwhile...

- Lots of protocols have started to use HTTP as a substrate.
- Other protocols want to use HTTP to bootstrap.
- Common requirement: “What URI should I use?”
- Well-Known URIs *seem* like they’re the answer.

- Well-Known URIs were not designed as a protocol bootstrap/tunnelling mechanism.
- Again, **limited** carve-out.
- Assuming that an origin corresponds to administrative boundaries can be problematic.
- Shoving everything into .well-known means that applications aren't using HTTP well (see: BCP56bis).
- There are very few cases where you can't just start with a URI (rather than a hostname).

1.1. Appropriate Use of Well-Known URIs

As per *[RFC7320]*, “publishing independent standards that mandate particular forms of URI substructure is inappropriate, because that essentially usurps ownership.” Well-known URIs are not an escape hatch from the requirements therein; they are a very limited carve-out of the path name space owned by the authority, ceded to standard use for a designated purpose.

That purpose is to facilitate discovery of information about an origin 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 the information applies to many (or all) of the origin’s resources.

Typically, the resource(s) identified by a well-known URI will make information about the origin (e.g., policy) available directly, or provide references to other URIs that provide it. In general, that information should be applicable to most origins (i.e., Web sites – while acknowledging that some origins might not use a particular well-known location, for various reasons).

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.

Specifically, well-known URIs are not a “protocol registry” for applications and protocols that wish to use HTTP as a substrate. Instead, such applications and protocols are encouraged to use an absolute URI to bootstrap their operation, rather than using a hostname and a well-known URI.

Exceptionally, the registry expert(s) may approve such a registration for documents in the IETF Stream [*RFC5741*], in consultation with the IESG, provided that the protocol in question cannot be bootstrapped with a URI (e.g., the discovery mechanism can only carry a hostname). However, merely making it easier to locate it is not a sufficient reason. Likewise, future use unsupported by the specification in question is not sufficient reason to register a well-known location.

Well-known locations are also not suited for information on topics other than the origin that they are located upon; for example, creating a well-known resource about a business entity or organisational structure presumes that Internet hosts and organisations share structure, and are likely to have significant deployment issues in environments where this is not true.

Next Steps

- Clarity on this would be very helpful, as we've had a number of late-stage “discussions” between ADs, registry expert and WGs about how to use well-known URIs
- RFC5785 was AD-sponsored
- AD-sponsor bis?