

5785bis

IETF101, London

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

“...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.”

-RFC5785

- 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 IETF protocols have started to use HTTP as a substrate.
- 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).

Proposal

- Expand upon “Appropriate Use” section
- This is a clarification; no new normative requirements
- 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?