

Workgroup: Thing-to-Thing Research Group

Internet-Draft:

draft-bormann-t2trg-deref-id-00

Published: 5 November 2022

Intended Status: Informational

Expires: 9 May 2023

Authors: C. Bormann

Universität Bremen TZI

The "dereferenceable identifier" pattern

Abstract

In a protocol or an application environment, it is often important to be able to create unambiguous identifiers for some meaning (concept or some entity).

Due to the simplicity of creating URIs, these have become popular for this purpose. Beyond the purpose of identifiers to be uniquely associated with a meaning, some of these URIs are in principle *dereferenceable*, so something can be placed that can be retrieved when encountering such a URI.

The present -00 version is a stub to draw some attention to the opportunity that this pattern would benefit from a common description, documenting its benefits and pitfalls, and some mitigations for the latter.

About This Document

This note is to be removed before publishing as an RFC.

Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-bormann-t2trg-deref-id/>.

Discussion of this document takes place on the t2trg Research Group mailing list (<mailto:t2trg@irtf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/t2trg/>. Subscribe at <https://www.ietf.org/mailman/listinfo/t2trg/>.

Source for this draft and an issue tracker can be found at <https://github.com/cabo/deref-id>.

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 9 May 2023.

Copyright Notice

Copyright (c) 2022 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/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.

Table of Contents

- [1. Introduction](#)
- [2. Examples for "dereferenceable identifiers"](#)
 - [2.1. Protocol and Protocol Version identifiers](#)
 - [2.2. Concept identifiers](#)
 - [2.3. MORE EXAMPLES](#)
- [3. Pitfalls](#)
 - [3.1. Server overload](#)
 - [3.2. Longevity of identifiers](#)
 - [3.3. Redirect ambiguities](#)
- [4. IANA Considerations](#)
- [5. Security considerations](#)
- [6. Informative References](#)
- [Acknowledgements](#)
- [Author's Address](#)

1. Introduction

(Please see abstract.)

2. Examples for "dereferenceable identifiers"

This section is intended to present a number of examples where dereferenceable identifiers are in use in a protocol, including existing discussion about constraints on their usage, the benefits claimed for this constrained usage, and remaining issues.

2.1. Protocol and Protocol Version identifiers

Many protocols based on XML or JSON include a protocol or protocol version identifier in the heading to a data item.

E.g., [\[JSO\]](#) defines a language for data models that contain an identifier to the language version in use, here <https://json-schema.org/draft/2020-12/schema>. The model that can be retrieved from this URI in turn contains further dereferenceable identifiers that point to further details.

[Section 8.1.1](#) of [\[JSO\]](#) has this:

If this URI identifies a retrievable resource, that resource SHOULD be of media type "application/schema+json".

So it acknowledges that the dereferenceability is optional, but does place further restrictions on what can be the result of a successful dereference: another one of these data models, which in turn contain further dereferenceable identifiers.

2.2. Concept identifiers

The *problem details* format [\[PROBLEM\]](#) uses a dereferenceable identifier for its "type" field. The value is a URI that "identifies the specific "problem type" (e.g., "out of credit")" ([Section 1](#) of [\[PROBLEM\]](#)).

[Section 3.1.1](#) of [\[PROBLEM\]](#) has this:

If the type URI is a locator (e.g., those with a "http" or "https" scheme), dereferencing it SHOULD provide human-readable documentation for the problem type (e.g., using HTML [\[HTML5\]](#)).

but then warns:

However, consumers SHOULD NOT automatically dereference the type URI, unless they do so when providing information to developers (e.g., when a debugging tool is in use).

[Section 5](#) of [\[PROBLEM\]](#) further details:

A problem's type URI SHOULD resolve to HTML [\[HTML5\]](#) documentation that explains how to resolve the problem.

This becomes even more interesting as [Section 5.2](#) of [\[PROBLEM\]](#) then gives this advice:

Registrations MAY use the prefix "https://iana.org/assignments/http-problem-types#" for the type URI.

A reference to the place where registrations for these items are managed is certainly desirable, however, the implications on the management of fragment identifiers in the HTML documents that IANA generates from registration information are an example for the increased complexity dereferenceable identifiers may place on the owners of the URI space employed.

2.3. MORE EXAMPLES

There are a lot more examples in published RFCs; add them to this document.

3. Pitfalls

3.1. Server overload

If a data item containing dereferenceable identifier(s) becomes widely distributed, naive implementations that handle such a data item might dereference these identifiers as part of a routine operation. Many definitions of dereferenceable identifiers contain admonitions that such a behavior can cause an implosion of requests on the server(s) for the URI.

3.2. Longevity of identifiers

Dereferenceable URIs usually contain domain names, whose ownership can change. As a result, and for other reasons as well, parts of the name space of an origin may come under new administration, which can change the policies that apply to resources made available there.

These are problems of such URIs in general (and can be mitigated by going to a non-dereferenceable kind of URIs such as one based on the 'tag' uri scheme [[TAG](#)]). However, the problems are exacerbated by their use as a dereferenceable identifier. The new owner/administrator might more easily accept that a certain chunk of their URI space should not be used (which suffices for a non-dereferenceable identifier based on this kind of URI namespace) than that certain content needs to be offered there (potentially presenting non-trivial loads, some mechanisms needed to update that information, and legal liabilities that are hard to assess).

3.3. Redirect ambiguities

Dereferencing an identifier may involve following some redirections; whether that following is actually implied, or desired (or even desirable) is rarely being discussed.

4. IANA Considerations

This document makes no concrete requests on IANA, but does point out that IANA resources might be a good target for a certain class of dereferenceable identifiers.

5. Security considerations

The ability to create a denial of service attack by pointing a dereferenceable identifier into a popular data item that is widely distributed is implied by the discussion in [Section 2](#), alongside with some recommendations for implementers that would mitigate such attacks. A problem with such recommendations is that they need to be followed by implementations that are using dereferenceable identifiers, which might not care much.

6. Informative References

- [JS0] Wright, A., Andrews, H., Hutton, B., and G. Dennis, "JSON Schema: A Media Type for Describing JSON Documents", Work in Progress, Internet-Draft, draft-bhutton-json-schema-01, 10 June 2022, <<https://www.ietf.org/archive/id/draft-bhutton-json-schema-01.txt>>.
- [PROBLEM] Nottingham, M., Wilde, E., and S. Dalal, "Problem Details for HTTP APIs", Work in Progress, Internet-Draft, draft-ietf-httpapi-rfc7807bis-04, 5 September 2022, <<https://www.ietf.org/archive/id/draft-ietf-httpapi-rfc7807bis-04.txt>>.
- [TAG] Kindberg, T. and S. Hawke, "The 'tag' URI Scheme", RFC 4151, DOI 10.17487/RFC4151, October 2005, <<https://www.rfc-editor.org/info/rfc4151>>.

Acknowledgements

Christian Amsüss pointed out that this document would be good to have.

Author's Address

Carsten Bormann
Universität Bremen TZI
Postfach 330440
D-28359 Bremen
Germany

Phone: [+49-421-218-63921](tel:+49-421-218-63921)
Email: cabo@tzi.org