

Network Working Group
Internet-Draft
Expires: June 20, 2003

D. Connolly
W3C
M. Baker
December 20, 2002

**A Registry of Assignments using Ubiquitous Technologies and Careful
Policies
draft-connolly-w3c-accessible-registries-00**

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on June 20, 2003.

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

Abstract

Registries are a critical part of many of the systems which inhabit the Internet today. Unfortunately, not a lot of the assignments within those registries are as accessible as they could be. This document outlines a best current practice approach to improving the accessibility of assignments within registries using ubiquitous technologies and low cost supporting policies.

1. Introduction

The success of the World Wide Web project has demonstrated the value of dereferencable URIs, and linking information together using them. Many registries on the Internet today create assignments whose principle identifier is not a URI, preventing the information associated with that assignment from being integrated with Web tools such as search engines, and from being related to other information, for example linking the assignment of the "application/xml" Internet media type to [RFC 3023](#).

On occasion, where an assignment has been provided a URI, it is not done so authoritatively, reducing the value of the URI as an identifier, as references to the assignment are split between the URI and the non-URI form. For example, the Internet media type "application/pdf" could also be identified by the URI "http://www.iana.org/assignments/media-types/application/pdf", such that new protocols could use the URI instead of "application/pdf". Anybody inspecting a message would be able to dereference the URI to discover that it was in fact a media type, rather than simply a string that happens to resemble one.

2. Proposed Solution

As the title of this note intimates, we suggest that a combination of the use of currently ubiquitous technologies, together with the implementation of careful policies, comprises a practical and long-term viable solution to this problem.

2.1 Ubiquitous Technologies

The basis of our proposed solution is for registries to use dereferencable URIs, in particular http: scheme URIs, as the principle identifier syntax for assignments. When dereferenced, these URIs should return at least some human processable information describing the assignment; the requestor, the details of the assignment request, URIs of relevant specifications, etc..

HTTP [[1](#)] is our recommended dereferencing mechanism, because it includes features specifically for managing the evolution of resources (of which registry assignments are one kind), and has been highly optimized for the transfer of coarse grained data objects. In particular;

redirection: permits the registry to communicate the movement of assignments, or, should it be required, the splitting of an assignment into multiple assignments

content negotiation: permits the URI of the assignment to be used to retrieve different data formats, should it be necessary to support more than one

metadata: information such as that provided by the "Last-Modified" header, used with the "If-Modified-Since" method qualifier permits efficient dereferencing in the face of mostly static data, such as most registry assignments

FTP is another protocol that could be used if required, but its lack of redirection and content negotiation makes URIs in the ftp: URI scheme more prone to change over time (and to create additional URIs, for example to support additional data formats), in addition to requiring an extra network round trip for authentication, even for anonymous transfers.

We also recommend that IANA and IETF use iana.org and ietf.org respectively in the authority component of their URIs. See [Appendix A](#) for a FAQ on this issue.

2.2 Careful Policies

2.2.1 Dereferenceable

Assignments, whether using the http: or ftp: URI scheme, should be dereferenceable (via HTTP GET or FTP RETR, respectively) to some information about the assignment. For example, the assignment for an Internet media type would ideally list information such as the RFC which includes the IANA required registration form, the date it was registered, etc..

This conclusion was made by the W3C Technical Architecture Group (TAG) in a finding entitled "Mapping between URIs and Internet Media Types" [2].

2.2.2 Persistence

The registry must commit to supporting and documenting a persistence policy for assignments. The policy need not be the same for all assignments, but should at least include the following;

- Moving assignments will be avoided at almost any cost

- Should moving be required;

 - A one year notice will be given and broadly announced

 - Redirection services (for http: scheme URIs) will be provided for three years after the move

3. Acknowledgements

The authors would like to thank the contributions of Larry Masinter, as well as the members of the W3C Technical Architecture Group.

In addition, Tim Berners-Lee's work on the World Wide Web project provided significant input to, and motivation for, this draft.

References

- [1] 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.
- [2] Williams, S., "Mapping between URIs and Internet Media Types", May 2002, <<http://www.w3.org/2001/tag/2002/01-uriMediaType-9>>.
- [3] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part One: The Comprehensive DDDS", [RFC 3401](#), October 2002.
- [4] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part Two: The Algorithm", [RFC 3402](#), October 2002.
- [5] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part Three: The Domain Name System (DNS) Database", [RFC 3403](#), October 2002.
- [6] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part Four: The Uniform Resource Identifiers (URI) Resolution Application", [RFC 3404](#), October 2002.
- [7] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part Five: URI.ARPA Assignment Procedures", [RFC 3405](#), October 2002.
- [8] Daigle, L., van Gulik, D., Iannella, R. and P. Falstrom, "Uniform Resource Names (URN) Namespace Definition Mechanisms", [RFC 3406](#), October 2002.

Authors' Addresses

Daniel W. Connolly
World Wide Web Consortium
MIT Laboratory for Computer Science
545 Technology Square
Cambridge, MA 02139
US

Phone: tel:+1-913-491-0501
EMail: <mailto:connolly@w3.org>
URI: <http://www.w3.org/People/Connolly/>

Mark Baker

Phone: tel:+1-613-286-4390

E-Mail: <mailto:distobj@acm.org>

URI: <http://www.markbaker.ca/>

[Appendix A](#). FAQs

[A.1](#) What if 'iana.org' is disputed and reassigned to somebody else?

It is extremely unlikely that this will happen. IANA owns its name, as ISOC is a legal entity. The risk of this happening is certainly lower than other relevant risks.

[A.2](#) Wouldn't URNs, or some other DNS-independant form of URI be a better solution?

We don't believe so, no. First, we feel that it is extremely important that assignments be dereferenceable using ubiquitous technologies so that information about the assignment can be easily accessed by anyone (which is not currently the case for URNs). Second, any registry of names (such as an URN namespace) that is, or becomes important, will eventually find itself facing the same issues (e.g. trademark) that DNS has faced. DNS, and URI schemes using it, are at least "the devil we know".

Finally, we feel that URNs do not address the issue of "URI persistence" as is claimed in [RFC 3404](#) [6], part of the DDDS[3][4][5][6][7][8] set of specifications. Specifically, the introduction of [RFC 3404](#) states;

Uniform Resource Identifiers (URI) have been a significant advance in retrieving Internet-accessible resources. However, their brittle nature over time has been recognized for several years. The Uniform Resource Identifier working group proposed the development of Uniform Resource Names (URN) to serve as persistent, location-independent identifiers for Internet resources in order to overcome most of the problems with URIs. [RFC 1737](#) sets forth requirements on URNs.

We believe that URNs are no less brittle than other URIs as persistent identifiers, and that "http://www.w3.org/" is no more location-dependant than, for example "urn:w3c", should it ever become resolvable. No URI presents a single point of failure to resolution, as the existence of HTTP caches demonstrate, since they can cache representations of any resource identified by any URI, not just those using the http: URI scheme.

[A.3](#) What if IANA's server goes down?

XSLT processors don't stop working when the W3C's Web server goes offline for a day (as it has in the past), despite referencing the w3.org XSLT namespace; they don't generally consult it in real-time. Any services that do depend on live access will have to have caching/

redundancy/failover/robustness mechanisms; the same mechanisms they would need to have should their own network go down.

Full Copyright Statement

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

