

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 20, 2014

D. Thaler
Microsoft
October 17, 2013

**Guidelines and Registration Procedures for New URI Schemes: Problem
Statement
draft-thaler-uri-scheme-reg-ps-01.txt**

Abstract

This document describes some problems with the existing guidelines and procedures, as documented in [RFC 4395](#), for new URI schemes.

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 <http://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 April 20, 2014.

Copyright Notice

Copyright (c) 2013 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
2.	Problems	4
2.1.	Current registration process doesn't scale well	4
2.2.	Lack of incentive to register	5
2.3.	Current private scheme guidance causes conflicts	5
3.	Security Considerations	6
4.	IANA Considerations	6
5.	Informative References	6
	Author's Address	7

[1.](#) Introduction

[RFC 4395](#) [[RFC4395](#)] provides guidelines and recommendations for the definition of Uniform Resource Identifier (URI) schemes. It defines procedures and guidelines for four types of URI schemes:

- a. Permanent, which [[RFC4395](#)] requires for all IETF Standards-Track schemes, and which has strict requirements.
- b. Provisional, which has a lower barrier.
- c. Historical, which is for schemes no longer in use and hence generally does not apply to "new" URI schemes.
- d. Private, meaning not registered with IANA.

As explained in [Section 1 of \[RFC4395\]](#), the purpose of an IANA-maintained registry is to:

1. provide a central point of discovery for established URI scheme names, and easy location of their defining documents;
2. discourage use of the same URI scheme name for different purposes;
3. help those proposing new URI scheme names to discern established trends and conventions, and avoid names that might be confused with existing ones;
4. encourage registration by setting a low barrier for provisional registrations.

However, the guidance in [[RFC4395](#)] is, in many cases that are now common, ambiguous or insufficient to accomplish the stated purposes. This document discusses a number of such problems. In doing so, we note that an effort was started to update the guidance, in

[[I-D.ietf-iri-4395bis-irireg](#)]. It does not, however, address the problems we discuss in this document, although it may be the logical place to do so.

It is first important to understand the scale of the problem. It is already common on many widely deployed platforms (including Windows, iOS, and Android) and form factors (PCs, phones, etc.) today to allow applications to be associated with specific URI schemes, such that when the URI is accessed (e.g., clicking on a link in a browser, or calling an equivalent API from an application), the associated application is launched to handle the URI. That is, the application is given the URI and determines what action to take (as opposed to being given content that the URI points to). Indeed, some such URIs are simply Uniform Resource Names that contain the data themselves, rather than Uniform Resource Locators that can be resolved to content. As such, URIs are increasingly becoming a form of inter-process communication as a way to invoke another application, with arguments placed in the scheme-specific part of the URI. Thus, in the extreme case, every application might define its own URI scheme, and the number of applications available on mainstream platforms today is easily numbered in the hundreds of thousands.

This use of URIs can be viewed as different from the web. That is, an increasingly larger portion of URI schemes are intended for "local" use, rather than for use with the web. The "URI Generic Syntax" [[RFC3986](#)] explicitly allows for such a wide scope of use of URIs. It states, in [section 1.1](#):

This specification does not limit the scope of what might be a resource; rather, the term "resource" is used in a general sense for whatever might be identified by a URI. Familiar examples include an electronic document, an image, a source of information with a consistent purpose (e.g., "today's weather report for Los Angeles"), a service (e.g., an HTTP-to-SMS gateway), and a collection of other resources. A resource is not necessarily accessible via the Internet; e.g., human beings, corporations, and bound books in a library can also be resources. Likewise, abstract concepts can be resources, such as the operators and operands of a mathematical equation, the types of a relationship (e.g., "parent" or "employee"), or numeric values (e.g., zero, one, and infinity).

and

This specification does not place any limits on the nature of a resource, the reasons why an application might seek to refer to a resource, or the kinds of systems that might use URIs for the sake of identifying resources.

The current process was designed based in part on joint recommendations from the W3C and IETF in 2002 [[RFC3305](#)], when the known uses of schemes were such that there were 34 registered schemes, 51 known publically documented but unregistered schemes, and 50 or so private schemes with 2-3 being added every day, as noted (see [Section 3.1 of \[RFC3305\]](#)). Such private growth has continued and expanded to more platforms since then, such that the public schemes are now probably a small minority.

2. Problems

2.1. Current registration process doesn't scale well

[Section 5.2 of \[RFC4395\]](#) requires a four-week mailing list review for all Permanent registrations. It is, however, ambiguous as to whether a mailing list review is required for Provisional registrations and if so, for how long. The longer the process, the less of an incentive there is to register Provisional schemes. This problem was discussed in 2010 by the IRI WG, which concluded that a mailing list review should not be required for Provisional schemes, only expert review which may take up to two weeks, but this conclusion has not yet been documented.

The manual step of expert review still introduces a scalability bottleneck. What if all new applications being submitted to an app store started sending requests for Provisional URI schemes? The expert review process would be overwhelmed, especially if no one is paid to do the expert review. As such, the goals stated in [Section 1](#) become far less effective when registered schemes are only a tiny fraction of the URI schemes in use in practice.

The author ran an experiment in 2012, which was reported to the IRI WG at its final meeting at IETF 85, where over 75 schemes that were listed on Wikipedia as being unregistered but in use were submitted as third-party registrations. All of them were registered after two weeks had passed and it was pointed out that the deadline had expired and per the process in [\[RFC4395\]](#), must be automatically listed. The only noticeable outcome of the expert review, other than to introduce a two week delay and manual effort, was to add a warning about the unknown security impact of one scheme. This is not intended to imply that the expert review was not valuable, only that the value provided could not scale effectively if the process were stressed with the current potential demand.

In summary, [\[RFC4395\]](#) defines a set of goals, which we listed above in [Section 1](#). The current mechanism does not meet those goals. To meet the stated goals would require the majority of schemes to be registered. The current process cannot scale to do so, given current

practice. Hence, we either need to change the goals, or change the process, or both.

2.2. Lack of incentive to register

Currently there is little incentive for an organization outside the IETF to register schemes (whether as Permanent, Provisional, or Historical). Registering introduces a cost, both in terms of manual effort needed to apply, but also in the time delay introduced. This cost must be weighed against the benefit, which is primarily to simply lower the risk of collision. (Another benefit is to provide ease of access to relevant documentation via the IANA registry, although this benefit is often seen as unimportant or even undesirable in some cases.)

As long as the risk of collision is perceived to be low, or the effect of collision considered to be acceptable (e.g., asking the user which app to launch), registration is bypassed in favor of a "Private" scheme. The effect of collision can of course be problematic (though the scheme-defining organization may not realize the danger) when the syntax of the scheme-specific part differs. Launching an application with a URI that is invalid according to that application's syntax for the custom URI scheme is not useful.

An app store certification process could in theory require or encourage Provisional application. However, there is little incentive for them to do so either, since an app store itself has a process which would be delayed and disincen application developers to submit applications.

2.3. Current private scheme guidance causes conflicts

[Section 2.8 of \[RFC4395\]](#) states:

Organizations that desire a private name space for URI scheme names are encouraged to use a prefix based on their domain name, expressed in reverse order. For example, a URI scheme name of com-example-info might be registered by the vendor that owns the example.com domain name.

There are multiple problems with the above guidance:

1. No guidance is given for when it might or might not be appropriate to use a private name space. For example, is this guidance appropriate for application vendors defining a custom scheme that they want to associate the application with? As such, the current assumption is that it is appropriate for anyone who can live with some potential risk of collision.

2. Hyphens occur in actual domain names. Consider one organization that owns the domain name "foo.bar.example", and another organization that owns "foo-bar.example". Using the mechanism implied in the example can result in both colliding with "example-bar-foo-info".
3. The guidance is only an encouragement, and no precise algorithm is given. For example, whether "." should be converted to "-" as in the example is unclear. If an organization is actually trying to follow the recommended guidelines, they will likely use a "-" as directed and risk conflicts as noted above. More commonly, an organization today will simply use a string that identifies (say) their application, and not be based on a domain name.
4. No protection is suggested against IANA later granting registration to a scheme that follows the recommended convention that is in use by someone else. For example, as can be seen at [\[IANAURI\]](#), there are already registered schemes that use "." (e.g., "iris.beep") and "-" (e.g., "xcon-userid") in them, and there could be similar new schemes registered at any time. If an organization had previously acquired the TLD "iris" or "xcon", those values could already be in use in applications from those organizations. Especially now that ICANN is allowing gTLD applications, this is a very real possibility.

3. Security Considerations

The security considerations in [\[RFC4395\]](#) still apply.

4. IANA Considerations

This document requires no actions by the IANA.

5. Informative References

[I-D.ietf-iri-4395bis-irireg]

Hansen, T., Hardie, T., and L. Masinter, "Guidelines and Registration Procedures for New URI/IRI Schemes", [draft-ietf-iri-4395bis-irireg-04](#) (work in progress), December 2011.

[IANAURI] IANA, ., "Uniform Resource Identifier (URI) Schemes", 2013, <<http://www.iana.org/assignments/uri-schemes/uri-schemes.xhtml>>.

[RFC3305] Mealling, M. and R. Denenberg, "Report from the Joint W3C/IETF URI Planning Interest Group: Uniform Resource Identifiers (URIs), URLs, and Uniform Resource Names

(URNs): Clarifications and Recommendations", [RFC 3305](#), August 2002.

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.

[RFC4395] Hansen, T., Hardie, T., and L. Masinter, "Guidelines and Registration Procedures for New URI Schemes", [BCP 35](#), [RFC 4395](#), February 2006.

Author's Address

Dave Thaler
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
USA

Phone: +1 425 703 8835
Email: dthaler@microsoft.com

