

URNBIS
Internet-Draft
Obsoletes: [3406](#) (if approved)
Intended status: Best Current Practice
Expires: August 16, 2014

P. Saint-Andre
&yet
February 12, 2014

Uniform Resource Name (URN) Namespace Definition Mechanisms
draft-ietf-urnbis-rfc3406bis-urn-ns-reg-09

Abstract

This document supplements the Uniform Resource Name (URN) syntax specification by defining the concept of a URN namespace, as well as mechanisms for defining and registering such namespaces. This document obsoletes [RFC 3406](#).

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 August 16, 2014.

Copyright Notice

Copyright (c) 2014 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.	Terminology	3
3.	What is a URN Namespace?	3
4.	URN Namespace Types	5
4.1.	Formal Namespaces	5
4.2.	Informal Namespaces	6
5.	Defining a URN Namespace	7
5.1.	Purpose	7
5.2.	Syntax	8
5.3.	Assignment	8
5.4.	Security and Privacy	9
5.5.	Resolution	9
6.	Registration Template	10
6.1.	Namespace ID	10
6.2.	Version	10
6.3.	Date	10
6.4.	Registrant	10
6.5.	Purpose	10
6.6.	Syntax	10
6.7.	Assignment	10
6.8.	Resolution	10
6.9.	Documentation	11
7.	Registering a URN Namespace	11
7.1.	Formal Namespaces	11
7.2.	Informal Namespaces	11
8.	Guidelines for Designated Experts	12
9.	IANA Considerations	12
10.	Security and Privacy Considerations	12
11.	References	12
11.1.	Normative References	12
11.2.	Informative References	13
Appendix A.	Changes from RFC 3406	14
Appendix B.	Contributors	14
Appendix C.	Acknowledgements	14
	Author's Address	14

[1.](#) Introduction

A Uniform Resource Name (URN) [[I-D.ietf-urnbis-rfc2141bis-urn](#)] is a Uniform Resource Identifier (URI) [[RFC3986](#)] that is intended to serve as a persistent, location-independent resource identifier. This document supplements the Uniform Resource Name (URN) syntax specification [[I-D.ietf-urnbis-rfc2141bis-urn](#)] by defining:

1. The concept of a URN namespace.

Saint-Andre

Expires August 16, 2014

[Page 2]

2. A mechanism for defining a URN namespace and associating it with a public identifier (called a Namespace ID or "NID").
3. Procedures for registering NIDs with the Internet Assigned Numbers Authority (IANA).

Syntactically, the NID follows the 'urn' scheme name. For instance, a URN in the 'example' namespace [[RFC6963](#)] might be of the form "urn:example:foo".

This document rests on two key assumptions:

1. Assignment of a URN is a managed process.

A string that conforms to the URN syntax is not necessarily a valid URN, because a URN needs to be assigned according to the rules of a particular namespace (in terms of syntax, semantics, and process).

2. The space of URN namespaces is itself managed.

A string in the namespace identifier slot of the URN syntax is not necessarily a valid URN namespace identifier, because in order to be valid a namespace needs to be defined and registered in accordance with the rules of this document.

URN namespaces were originally defined in [[RFC2611](#)], which was obsoleted by [[RFC3406](#)]. Based on experience with defining and registering URN namespaces since that time, this document specifies URN namespaces with the smallest reasonable set of changes from [[RFC3406](#)], while at the same time simplifying the registration process. This document obsoletes [RFC 3406](#).

2. Terminology

Several important terms used in this document are defined in the URN syntax specification [[I-D.ietf-urnbis-rfc2141bis-urn](#)].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

3. What is a URN Namespace?

A URN namespace is a collection of identifiers that are (1) unique, (2) assigned in a consistent way, and (3) assigned according to a common definition.

1. The "uniqueness" constraint means that an identifier within the namespace is never assigned to more than one resource and never reassigned to a different resource, even if the identifier itself is deprecated or becomes obsolete.
2. The "consistent assignment" constraint means that an identifier within the namespace is assigned by an organization or created in accordance with a process or algorithm that is always followed.
3. The "common definition" constraint means that there are clear definitions for the syntax of identifiers within the namespace and the process of assigning or creating them.

A URN namespace is identified by a particular NID in order to ensure the global uniqueness of URNs and, optionally, to provide a cue regarding the structure of URNs assigned within a namespace.

With regard to global uniqueness, using different NIDs for different collections of identifiers ensures that no two URNs will be the same for different resources, since each collection is required to uniquely assign each identifier. However, a single resource can have more than one URN assigned to it for different purposes (e.g., some numbers might be valid identifiers in two different identifier systems, where the namespace identifier differentiates between the resulting URNs).

With regard to the structure of URNs assigned within a namespace, the development of an identifier structure (and thereby a collection of identifiers) depends on the requirements of the community defining the identifiers, how the identifiers will be assigned and used, etc. These issues are beyond the scope of URN syntax and the general rules for URN namespaces, because they are specific to the community defining a namespace (e.g., the bibliographic and publishing communities in the case of the 'ISBN' and 'ISSN' namespaces, or the developers of extensions to the Extensible Messaging and Presence Protocol in the case of the 'XMPP' namespace).

URN namespaces inherit certain rights and responsibilities by the nature of URNs [[I-D.ietf-urnbis-rfc2141bis-urn](#)], e.g.:

1. They uphold the general principles of a well-managed URN namespace by providing persistent identification of resources and unique assignment of identifier strings.
2. They can be registered in global registration services.

4. URN Namespace Types

There are two types of URN namespace: formal and informal. These are distinguished by the expected level of service, the information needed to define the namespace, and the procedures for registration. Because the majority of the namespaces registered so far have been formal, this document concentrates on formal namespaces.

Note: [[RFC3406](#)] defined a third type of "experimental namespaces", denoted by prefixing the namespace identifier with the string "X-". Consistent with [[RFC6648](#)], this specification removes the experimental category. Because experimental namespaces were never registered, removing the experimental category has no impact on the existing registries or future registration procedures.

4.1. Formal Namespaces

A formal namespace provides benefit to some subset of users on the Internet (e.g., it would not make sense for a formal namespace to be used only by a community or network that is not connected to the Internet). For example, it would be inappropriate for a NID to effectively force someone to use a proprietary network or service not open to the general Internet user. The intent is that, while the community of those who might actively use the names assigned within that NID might be small, the potential use of identifiers within that NID is open to any user on the Internet. Formal NIDs might be appropriate when some aspects are not fully open. For example, a namespace might make use of a fee-based, privately managed, or proprietary registry for assignment of URNs in the namespace. However, it might still benefit some Internet users if the associated services have openly-published access protocols.

An organization that will assign URNs within a formal namespace ought to meet the following criteria:

1. Organizational stability and the ability to maintain the URN namespace for a long time; absent such evidence, it ought to be clear how the namespace can remain viable if the organization can no longer maintain the namespace.
2. Competency in name assignment. This will improve the likelihood of persistence (e.g. to minimize the likelihood of conflicts).
3. Commitment to not reassigning existing names and to allowing old names to continue to be valid, even if the owners or assignees of those names are no longer members or customers of that organization. With regard to URN resolution [[RFC2276](#)], this does

not mean that there needs to be resolution of such names, only that the names will not resolve to false or stale information.

A formal namespace establishes a particular NID, subject to the following constraints (above and beyond the syntax rules specified in [[I-D.ietf-urnbis-rfc2141bis-urn](#)]):

1. It MUST NOT be an already-registered NID.
2. It MUST NOT start with "urn-" (which is reserved for informal namespaces).
3. It MUST be more than two characters long.
4. It MUST NOT start with "aa-", where "aa" is any combination of two ASCII letters.
5. It MUST NOT start with the string "xn--", which is reserved for potential representation of DNS A-labels in the future [[RFC5890](#)].

All two-letter combinations, and all two-letter combinations followed by "-" and any sequence of valid NID characters, are reserved for potential use as countrycode-based NIDs for eventual national registrations of URN namespaces. The definition and scoping of rules for allocation of responsibility for such countrycode-based namespaces is beyond the scope of this document.

[4.2.](#) Informal Namespaces

Informal namespaces are full-fledged URN namespaces, with all the associated rights and responsibilities. Informal namespaces differ from formal namespaces in the process for assigning a NID: for an informal namespace, the registrant does not designate the NID; instead, IANA assigns a NID consisting of the string 'urn-' followed by one or more digits (e.g., "urn-7") where the digits consist of the next available number in the sequence of positive integers assigned to informal namespaces. Thus the syntax of an informal namespace is:

"urn-" <number>

The only restrictions on <number> are that it (1) consist strictly of ASCII digits and (2) not cause the NID to exceed the length limitations defined in the URN syntax specification [[I-D.ietf-urnbis-rfc2141bis-urn](#)].

[5.](#) Defining a URN Namespace

The definition of a formal namespace ought to pay particular attention to:

1. The purpose of the namespace.
2. The syntax of URNs assigned within the namespace.
3. The process for assigning URNs within the namespace.
4. The security implications of assigning and using the assigned URNs.
5. Optionally, the process for resolving URNs issued within the namespace.

The following sections explain these matters in greater detail. For convenience, a template for defining and registering a URN namespace is provided under [Section 6](#). This information can be especially helpful to entities that wish to request assignment of a URN in a namespace and to entities that wish to provide URN resolution for a namespace.

[5.1.](#) Purpose

The "Purpose" section of the template describes matters such as:

1. The kinds of resources identified by URNs assigned within the namespace.
2. Why it is preferable to use URNs rather than some other technology (e.g., URIs) and why no existing URN namespace is a good fit.
3. The kinds of software applications that can use or resolve the assigned URNs (e.g., by differentiating among disparate namespaces, identifying resources in a persistent fashion, or meaningfully resolving and accessing services associated with the namespace).
4. The scope of the namespace (public vs. private, global vs. local to a particular organization, nation, or industry). For example, a namespace claiming to deal in "national identification numbers" ought to have a global scope and address all identity number structures, whereas a URN scheme for a particular national identification number system would need to handle only the structure for that nation's identity numbers.

5. How the intended community (and the Internet community at large) will benefit from using or resolving the assigned URNs.

5.2. Syntax

The "Syntax" section of the template describes:

1. A description of the structure of URNs within the namespace, in conformance with the fundamental URN syntax [[I-D.ietf-urnbis-rfc2141bis-urn](#)]. The structure might be described in terms of a formal definition (e.g., using Augmented BNF for Syntax Specifications (ABNF) as specified in [[RFC5234](#)]), an algorithm for generating conformant URNs, a regular expression for parsing the identifier into components, or by explaining that the structure is opaque.
2. Any special character encoding rules for assigned URNs (e.g., which character ought to always be used for single-quotes).
3. Rules for determining equivalence between two identifiers in the namespace. Such rules ought to always have the effect of eliminating false negatives that might otherwise result from comparison. If it is appropriate and helpful, reference can be made to the equivalence rules defined in the URI specification [[RFC3986](#)]. Examples of equivalence rules include equivalence between uppercase and lowercase characters in the Namespace Specific String, between hyphenated and non-hyphenated groupings in the identifier string, or between single-quotes and double-quotes. (Note that these are not normative statements for any kind of best practice related to handling of equivalences between characters in general; they are statements limited to one particular namespace only.)
4. Any special considerations necessary for conforming with the URN syntax. This is particularly applicable in the case of legacy naming systems that are used in the context of URNs. For example, if a namespace is used in contexts other than URNs, it might make use of characters that are reserved in the URN syntax. This section ought to note any such characters, and outline necessary mappings to conform to URN syntax. Normally, this will be handled by percent-encoding the character as specified in the URI specification [[RFC3986](#)].

5.3. Assignment

The "Assignment" section of the template describes matters such as:

1. Mechanisms or authorities for assigning URNs to resources. It ought to make clear whether assignment is completely open (e.g., following a particular algorithm), completely closed (e.g., for a private organization), or limited in various ways (e.g., delegated to authorities recognized by a particular organization); if limited, it ought to explain how to become an assigner of identifiers or how to request assignment of identifiers from existing assignment authorities.
2. Methods for ensuring that URNs within the namespace are unique. For example, identifiers might be assigned sequentially or in accordance with some well-defined process by a single authority, assignment might be partitioned among delegated authorities that are individually responsible for respecting uniqueness rules, or URNs might be created independently following an algorithm that itself guarantees uniqueness.

5.4. Security and Privacy

The "Security" section of the template describes any potential issues related to security and privacy with regard to assignment, use, and resolution of identifiers within the namespace. Examples of such issues include the consequences of producing false negatives and false positives during comparison for equivalence (see also [\[RFC6943\]](#)), leakage of private information when identifiers are communicated on the public Internet, the potential for directory harvesting, and various issues discussed in the guidelines for security considerations in RFCs [\[RFC3552\]](#) and the privacy considerations for Internet protocols [\[RFC6973\]](#).

5.5. Resolution

The "Resolution" section specifies the rules for resolution of URNs assigned within the namespace. If such URNs are intended to be resolvable, the namespace needs to be registered in a Resolution Discovery System (RDS, see [\[RFC2276\]](#)) such as DDDS. Resolution then proceeds according to standard URI resolution processes, as well as the mechanisms of the RDS. This section ought to list the requirements for becoming a recognized resolver of URNs in the relevant namespace (and being so listed in the RDS registry). Answers might include, but are not limited to:

1. The namespace is not listed with an RDS; therefore this section is not applicable.
2. Resolution mirroring is completely open, with a mechanism for updating an appropriate RDS.

3. Resolution is controlled by entities to which assignment has been delegated.

[6.](#) Registration Template

[6.1.](#) Namespace ID

Requested of IANA (formal) or assigned by IANA (informal).

[6.2.](#) Version

The version of the registration, starting with 1 and incrementing by 1 with each new version.

[6.3.](#) Date

The date when the registration is requested of IANA, using the format YYYY-MM-DD.

[6.4.](#) Registrant

The person or organization that has registered the NID, including the following information:

- o The name and address of the registering organization.
- o The name and contact information (email, phone number, and/or postal address) of the designated contact person.

[6.5.](#) Purpose

Described under [Section 5.1](#) of this document.

[6.6.](#) Syntax

Described under [Section 5.2](#) of this document.

[6.7.](#) Assignment

Described under [Section 5.3](#) of this document.

[6.8.](#) Resolution

Described under [Section 5.5](#) of this document.

6.9. Documentation

A pointer to an RFC, a specification published by another standards development organization, or another stable document that provides further information about the namespace.

7. Registering a URN Namespace

7.1. Formal Namespaces

The registration policy for formal namespaces is Expert Review as defined in the "IANA Considerations" document [[RFC5226](#)]. The key steps for registration of a formal namespace are:

1. Fill out the namespace registration template (see [Section 6](#)). This can be done as part of an Internet-Draft or a specification in another series, although that is not necessary.
2. Send the completed template to the urn-nid@ietf.org discussion list for review.
3. If necessary to address comments received, repeat steps 1 and 2.
4. If the designated experts approve the request, the IANA will register the requested NID.

A formal namespace registration can be revised by updating the registration template, following the same steps outlined above for new registrations. A revised registration should making special note of any relevant changes in the underlying technologies or namespace management processes.

7.2. Informal Namespaces

The registration policy for informal namespaces is First Come First Served [[RFC5226](#)]. The key steps for registration of an informal namespace are:

1. Write a completed namespace definition template (see [Section 6](#)).
2. Send it to the urn-nid@ietf.org discussion list for feedback.
3. Once the review period has expired, send the final template to IANA (via the iana@iana.org email address).

An informal namespace registration can be revised by updating the registration template, following the same steps outlined above for new registrations.

8. Guidelines for Designated Experts

Experience to date with NID registration requests has shown that registrants sometimes do not initially understand some of the subtleties of URN namespaces, and that defining the namespace in the form of a specification enables the registrants to clearly formulate their "contract" with the intended user community. Therefore, although the registration policy for formal namespaces is Expert Review and a stable specification is not strictly required, the designated experts for NID registration requests ought to encourage applicants to provide a stable specification documenting the namespace definition.

Naming can be difficult and contentious; the designated experts and applicants are strongly encouraged to work together in a spirit of good faith and mutual understanding to achieve rough consensus on progressing registrations through the process. They are also encouraged to bring additional expertise into the discussion if that would be helpful in adding perspective or otherwise resolving issues.

9. IANA Considerations

This document outlines the processes for registering URN namespaces, and has implications for the IANA in terms of registries to be maintained. In all cases, the IANA ought to assign the appropriate NID (formal or informal) once the procedures outlined in this document have been completed.

10. Security and Privacy Considerations

This document largely focuses on providing mechanisms for the declaration of public information. Nominally, these declarations will be of relatively low security profile, however there is always the danger of "spoofing" and providing misinformation. Information in these declarations ought to be taken as advisory.

The definition of a URN namespace needs to account for potential security and privacy issues related to assignment, use, and resolution of identifiers within the namespace.

11. References

11.1. Normative References

[I-D.ietf-urnbis-rfc2141bis-urn]
Saint-Andre, P., Ed., "Uniform Resource Name (URN) Syntax", [draft-ietf-urnbis-rfc2141bis-urn-07](#) (work in progress), January 2014.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.

11.2. Informative References

- [RFC2276] Sollins, K., "Architectural Principles of Uniform Resource Name Resolution", [RFC 2276](#), January 1998.
- [RFC2611] Daigle, L., van Gulik, D., Iannella, R., and P. Faltstrom, "URN Namespace Definition Mechanisms", [BCP 33](#), [RFC 2611](#), June 1999.
- [RFC3406] Daigle, L., van Gulik, D., Iannella, R., and P. Faltstrom, "Uniform Resource Names (URN) Namespace Definition Mechanisms", [BCP 66](#), [RFC 3406](#), October 2002.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", [BCP 72](#), [RFC 3552](#), July 2003.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), January 2008.
- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", [RFC 5890](#), August 2010.
- [RFC6648] Saint-Andre, P., Crocker, D., and M. Nottingham, "Deprecating the "X-" Prefix and Similar Constructs in Application Protocols", [BCP 178](#), [RFC 6648](#), June 2012.
- [RFC6943] Thaler, D., "Issues in Identifier Comparison for Security Purposes", [RFC 6943](#), May 2013.
- [RFC6963] Saint-Andre, P., "A Uniform Resource Name (URN) Namespace for Examples", [BCP 183](#), [RFC 6963](#), May 2013.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", [RFC 6973](#), July 2013.

[Appendix A](#). Changes from [RFC 3406](#)

This document makes the following substantive changes from [[RFC3406](#)]:

1. Relaxes the registration policy for formal namespaces from "IETF Review" to "Expert Review" [[RFC5226](#)].
2. Removes the category of experimental namespaces, consistent with [[RFC6648](#)].
3. Simplifies the registration template.

In addition, some of the text has been updated to be consistent with the definition of Uniform Resource Identifiers (URIs) [[RFC3986](#)] and the processes for registering information with the IANA [[RFC5226](#)], as well as more modern guidance with regard to security issues [[RFC3552](#)] and identifier comparison [[RFC6943](#)].

[Appendix B](#). Contributors

[RFC 3406](#), which provided the basis for this document, was authored by Leslie Daigle, Dirk-Willem van Gulik, Renato Iannella, and Patrik Faltstrom. Their work is gratefully acknowledged.

[Appendix C](#). Acknowledgements

Thanks to Marc Blanchet, Juha Hakala, Paul Jones, John Klensin, and Barry Leiba for their input.

Author's Address

Peter Saint-Andre
&yet

Email: ietf@stpeter.im

