

URNBIS
Internet-Draft
Obsoletes: [3406](#) (if approved)
Intended status: BCP
Expires: May 14, 2013

P. Saint-Andre, Ed.
Cisco Systems, Inc.
L. Daigle
D. van Gulik
R. Iannella
P. Faltstrom
November 10, 2012

Uniform Resource Name (URN) Namespace Definitions
draft-saintandre-urnbis-3406bis-00

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 May 14, 2013.

Copyright Notice

Copyright (c) 2012 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	3
2.	Discussion Venue	4
3.	Terminology	4
4.	What is a URN Namespace?	4
5.	URN Namespace Types	5
5.1.	Formal Namespaces	5
5.2.	Informal Namespaces	6
6.	Defining a URN Namespace	7
6.1.	Formal Namespaces	7
6.1.1.	Syntax	7
6.1.2.	Specification	8
6.2.	Informal Namespaces	8
7.	Registering a URN Namespace	9
7.1.	Formal Namespaces	9
7.2.	Informal Namespaces	9
8.	URN Namespace Definition Template	10
9.	Security Considerations	14
10.	IANA Considerations	14
11.	References	15
11.1.	Normative References	15
11.2.	Informative References	15
Appendix A.	Changes from RFC 3406	15
	Authors' Addresses	15

1. Introduction

A Uniform Resource Name (URN) is a Uniform Resource Identifier (URI) that is intended to serve as a persistent, location-independent resource identifier. The general class of URNs is differentiated from all other URIs through the use of the 'urn' URI scheme.

This document supplements the Uniform Resource Name (URN) syntax specification by defining (1) the concept of a URN namespace, (2) a mechanism for defining such namespaces and associating each namespace with a public identifier (called a Namespace ID or "NID"), and (3) procedures for registering such namespaces with the Internet Assigned Numbers Authority (IANA).

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, the goal of this document is to specify URN namespaces with the smallest reasonable set of changes from [[RFC3406](#)].

Although on the surface it might appear that this document is significantly different from [[RFC3406](#)], in general it only modifies the order of presentation, with the intent of making it easier for people to define and register URN namespaces. However, the only major substantive change is removing the category of experimental namespaces, in accordance with [[RFC6648](#)].

If approved, this document will obsolete [RFC 3406](#).

2. Discussion Venue

The discussion venue for this document is mailing list of the URNBIS WG; visit <<https://www.ietf.org/mailman/listinfo/urn>> for subscription and archive information.

3. Terminology

Several important terms used in this document are defined in the URN syntax specification [[I-D.saintandre-urnbis-2141bis](#)].

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](#)].

4. What is a URN Namespace?

For the purposes of URNs, a "namespace" is a collection of unique identifiers that are consistently assigned according to a common definition.

The uniqueness constraint means that an identifier within the namespace is never assigned to more than one resource and never re-assigned to a different resource (however, a single resource can have more than one URN assigned to it for different purposes).

The consistent assignment constraint means that an identifier within the namespace is assigned by an organization or in accordance with a process that is always followed.

The common definition constraint means that the syntax for identifiers within the namespace and the process for assigning such identifiers are clearly defined in a specification.

A URN namespace is identified by a particular designator (which syntactically follows the 'urn' scheme name) in order to:

- o Ensure the global uniqueness of URNs.
- o Optionally provide a cue regarding the structure of URNs assigned within a namespace.

With regard to global uniqueness, using different designators 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). For instance, some identifier

systems use strings of numbers as identifiers (e.g., ISBN, ISSN, phone numbers). It is conceivable that there might be some numbers that are valid identifiers in two different established identifier systems, and 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, is a process that is inherently dependent on the requirements of the community defining the identifier, how they will be assigned, and the uses to which they will be put. All of these issues are specific to the individual community seeking to define a namespace (e.g., a publishing community, an association of booksellers, developers of particular application protocols, etc.); therefore these issues are beyond the scope of URN syntax and the rules regarding URN namespaces in general.

URN namespaces inherit certain rights and responsibilities, including:

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

5. URN Namespace Types

There are two types of URN namespace: formal and informal. These are distinguished by the expected level of service, the information required to define the namespace, and the procedures for registration. To date, the vast majority of the registered namespaces have been formal, so 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.

5.1. Formal Namespaces

A formal namespace can be requested, and IETF review sought, in cases where the publication of the NID proposal and the underlying namespace will provide benefit to some subset of users on the Internet. That is, a formal NID proposal, if accepted, needs to be functional on and with the global Internet, not limited to users in communities or networks not connected to the Internet. For example,

consider a NID that is meant for naming of physics research; if that NID request required that the user use a proprietary network or service that was not at all open to the general Internet user, then it would make a poor request for a formal NID. The intent is that, while the community of those who may actively use the names assigned within that NID may be small (but no less important), the potential use of names within that NID is open to any user on the Internet.

It is expected that formal NIDs may be applied to namespaces where 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 may still provide benefit to some Internet users if the services associated have openly- published access protocols.

In addition to the basic information specified in the namespace definition template (see [Section 8](#)), a formal namespace request needs to be accompanied by documented considerations of the need for a new namespace and of the community benefit from formally establishing the proposed URN namespace.

Additionally, since the goal of URNs is to provide persistent identification, a formal namespace request needs to give some consideration as to the longevity and maintainability of the namespace. Possible factors to consider with regard to an organization that will assign URNs within a namespace include the following:

- o It ought to demonstrate 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.
- o It ought to demonstrate competency in name assignment. This will improve the likelihood of persistence (e.g. to minimize the likelihood of conflicts);
- o It ought to commit to not re-assigning 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, 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.

5.2. Informal Namespaces

Informal namespaces are full-fledged URN namespaces, with all the rights and responsibilities associated thereto. Informal namespaces differ from formal namespaces in the process for assigning a NID: IANA will assign an alphanumeric NID (e.g., "urn-7") to informal

namespaces, per the process outlined under [Section 7](#).

6. Defining a URN Namespace

A URN namespace is defined by the following factors:

- o The syntax of URNs assigned within the namespace, in conformance with the fundamental URN syntax [[I-D.saintandre-urnbis-2141bis](#)].
- o The process for assigning URNs within the namespace.
- o Optionally, the process for resolving URNs issued within the namespace.

Processes for resolution of URNs assigned within a namespace (if any) are out of scope for this document. The following sections provide guidelines for (1) defining the syntax of URNs within a namespace and (2) specifying how URNs will be assigned within a namespace.

6.1. Formal Namespaces

Formal NIDs are assigned as a result of IETF Consensus as defined in the "IANA Considerations" document [[RFC5226](#)]. Thus an application for a formal NID is made by publishing an RFC through normal IETF processes. The RFC need not be standards track (indeed, to date most RFCs registering URN namespaces have been informational), but it will be subject to IESG review and acceptance pursuant to the guidelines written here (as well as standard RFC publication guidelines).

6.1.1. Syntax

A formal namespace registration requests a particular NID, subject to the following constraints:

- o It MUST NOT be an already-registered NID.
- o It MUST NOT start with "urn-" (which is reserved for informal namespaces).
- o It MUST be more than 2 letters long.
- o It MUST NOT start with "XY-", where XY is any combination of two ASCII letters.

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.

6.1.2. Specification

The specification defining a formal namespace MUST include a completed namespace definition template (see [Section 8](#)).

The specification also MUST include the following sections.

The "Namespace Considerations" section outlines the perceived need for a new namespace (i.e., where existing namespaces fall short of the proposer's requirements). Potential considerations include:

- o Procedures for assigning URNs within this namespace
- o Processes for resolving URNs assigned within this namespace, if any
- o The type of resources to be identified
- o The type of services to be supported

It is expected that more than one namespace may serve the same "functional" purpose; the intent of the "Namespace Considerations" section is to provide a record of the proposer's "due diligence" in exploring existing possibilities, for the IESG's consideration.

The "Community Considerations" section explains how the intended community will benefit by publication of this namespace as well as how a general Internet user will be able to use the space if they care to do so. Potential considerations include:

- o Open assignment and use of identifiers within the namespace
- o Open operation of resolution servers for the namespace (server)
- o Creation of software that can meaningfully resolve and access services for the namespace (client)

The "IANA Considerations" section indicating that the document includes a URN NID registration that is to be entered into the IANA registry of URN NIDs.

6.2. Informal Namespaces

Informal namespaces are directly requested of IANA.

The namespace identifier assigned by IANA is a number sequence in the format:

"urn-" <number>

where <number> is chosen by the IANA on a First Come First Served basis as specified in the "IANA Considerations" document [[RFC5226](#)].

The only restrictions on <number> are (1) that it consist strictly of digits and (2) that it not cause the NID to exceed the length limitations defined in the URN syntax specification [[I-D.saintandre-urnbis-2141bis](#)].

7. Registering a URN Namespace

7.1. Formal Namespaces

The key steps for registration of a formal namespace are:

1. Write an Internet-Draft that includes all of the information described under Section FIXME.
2. Send the completed namespace definition template to the urn-nid@ietf.org discussion list for technical review.
3. Update the Internet-Draft as needed to address feedback, and repeat steps 2 and 3 as needed.
4. Based on comments received, update the Internet-Draft and repeat steps 2 and 3 as necessary.
5. Send a request to the IESG to publish the I-D as an RFC. The IESG may request further changes (published as I-D revisions) and/or direct discussion to designated working groups, area experts, etc.
6. If the IESG approves the document for publication as an RFC, send a request to IANA to register the requested NID.

A registration can be revised by updating the RFC through normal IETF processes [[RFC2606](#)]. The authors of the revised document need to follow the same steps outlined above for new registrations.

7.2. Informal Namespaces

The key steps for registration of an informal namespace are:

1. Complete the namespace definition template (see [Section 8](#)). This can be done as part of an Internet-Draft.
2. Send the completed template to the urn-nid@ietf.org discussion list for technical review.
3. Based on comments received, update the template and repeat steps 2 and 3 as necessary.
4. Once comments have been addressed and the review period has expired, send a registration request to IANA (via the iana@iana.org email address) with the final template.

Informal namespaces can also be revised by updating the template and processing it as outlined above for new registrations.

8. URN Namespace Definition Template

Definition of a URN namespace is accomplished by completing the following template. Apart from providing a mechanism for defining the structure of URNs assigned within the namespace, this information is designed to be useful for:

- o entities seeking to have a URN assigned in a namespace (if applicable)
- o entities seeking to provide URN resolvers for a namespace (if applicable)

This is particularly important for communities evaluating the possibility of using a portion of an existing URN namespace rather than creating their own.

As described under [Section 7.1](#), applications for formal URN namespaces MUST also document "Namespace Considerations", "Community Considerations" and "IANA Considerations".

The information to be provided in the template is as follows:

Namespace ID:

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

Registration Information:

This is information to identify the particular version of registration information:

- registration version number: starting with 1, incrementing by 1 with each new version
- registration date: date submitted to the IANA, using the format YYYY-MM-DD

Declared registrant of the namespace:

This includes:

Registering organization

Name

Address

Designated contact person

Name

Coordinates (at least one of email address, phone number, postal address)

Declaration of syntactic structure:

This section ought to outline any structural features of identifiers in this namespace. At the very least, this description may be used to introduce terminology used in other sections. This structure may also be used for determining realistic caching/shortcuts approaches; suitable caveats ought to be provided. If there are any specific character encoding rules (e.g., which character ought to always be used for single-quotes), these ought to be listed here.

Answers might include, but are not limited to:

- the structure is opaque (no exposition)
- a regular expression for parsing the identifier into components, including naming authorities

Relevant ancillary documentation:

This section ought to list any RFCs, standards, or other published documentation that defines or explains all or part of the namespace structure.

Answers might include, but are not limited to:

- RFCs outlining syntax of the namespace
- Other of the defining community's (e.g., ISO) documents outlining syntax of the identifiers in the namespace
- Explanatory material introducing the namespace

Identifier uniqueness considerations:

This section ought to address the requirement that URNs are assigned uniquely -- i.e., they are assigned to at most one resource, and are not reassigned.

(Note that the definition of "resource" is fairly broad; for example, information on "Today's Weather" might be considered a single resource, although the content is dynamic.)

Possible answers include, but are not limited to:

- exposition of the structure of the identifiers, and partitioning of the space of identifiers amongst assignment authorities which are individually responsible for respecting uniqueness rules
- identifiers are assigned sequentially
- information is withheld; the namespace is opaque

Identifier persistence considerations:

Although non-reassignment of URN identifiers ensures that a URN will persist in identifying a particular resource even after the "lifetime of the resource", some consideration ought to be given to the persistence of the usability of the URN. This is particularly important in the case of URN namespaces providing global resolution.

Possible answers include, but are not limited to:

- quality of service considerations

Process of identifier assignment:

This section ought to detail the mechanisms and/or authorities for assigning URNs to resources. It ought to make clear whether assignment is completely open, or if limited, how to become an assigner of identifiers, and/or get one assigned by existing assignment authorities.

Answers could include, but are not limited to:

- assignment is completely open, following a particular algorithm
- assignment is delegated to authorities recognized by a particular organization (e.g., the Digital Object Identifier Foundation controls the DOI assignment space and its delegation)
- assignment is completely closed (e.g., for a private organization)

Process for identifier resolution:

If a namespace is intended to be accessible for global resolution, it needs to be registered in an RDS (Resolution Discovery System, see [RFC 2276](#)) such as DDDS. Resolution then proceeds according to standard URI resolution processes, and the mechanisms of the RDS. What this section ought to outline is the requirements for becoming a recognized resolver of URNs in this namespace (and being so- listed in the RDS registry).

Answers may include, but are not limited to:

- the namespace is not listed with an RDS; therefore this section is not applicable
- resolution mirroring is completely open, with a mechanism

- for updating an appropriate RDS
- resolution is controlled by entities to which assignment has been delegated

Rules for Lexical Equivalence:

If there are particular algorithms for determining equivalence between two identifiers in the underlying namespace (hence, in the URN string itself), rules can be provided here.

Some examples include:

- equivalence between hyphenated and non-hyphenated groupings in the identifier string
- equivalence between single-quotes and double-quotes
- Namespace-defined equivalences between specific characters, such as "character X with or without diacritic marks".

Note that these are not normative statements for any kind of best practice for handling equivalences between characters; they are statements limited to reflecting the namespace's own rules.

Conformance with URN Syntax:

This section ought to outline any special considerations required 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 may make use of characters that are reserved in the URN syntax.

This section ought to flag any such characters, and outline necessary mappings to conform to URN syntax. Normally, this will be handled by hex encoding the symbol.

For example, see the section on SICIs in [RFC 2288](#).

Validation mechanism:

Apart from attempting resolution of a URN, a URN namespace may provide mechanisms for "validating" a URN -- i.e., determining whether a given string is currently a validly-assigned URN. There are two issues here: 1) users ought not "guess" URNs in a namespace; 2) when the URN namespace is based on an existing identifier system, it may not be the case that all the existing

identifiers are assigned on Day 0. The reasonable expectation is that the resource associated with each resulting URN is somehow related to the thing identified by the original identifier system, but those resources may not exist for each original identifier. For example, even if a URN namespace were defined based on telephone numbers, it is not clear that all telephone numbers would immediately become "valid" URNs, that could be resolved using whatever mechanisms are described as part of the namespace registration.

Validation mechanisms might be:

- a syntax grammar
- an on-line service
- an off-line service

Scope:

This section ought to outline the scope of the use of the identifiers in this namespace. Apart from considerations of private vs. public namespaces, this section is critical in evaluating the applicability of a requested NID. For example, a namespace claiming to deal in "social security numbers" ought to have a global scope and address all social security number structures (unlikely). On the other hand, at a national level, it is reasonable to propose a URN namespace for "this nation's social security numbers".

9. Security 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 mis-information. Information in these declarations ought to be taken as advisory.

10. 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 (informal or formal), as described above, once an IESG-designated expert has confirmed that the requisite registration process steps have been completed.

11. References

11.1. Normative References

- [I-D.saintandre-urnbis-2141bis]
Saint-Andre, P. and R. Moats, "Uniform Resource Name (URN) Syntax", [draft-saintandre-urnbis-2141bis-00](#) (work in progress), October 2012.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [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

- [RFC2606] Eastlake, D. and A. Panitz, "Reserved Top Level DNS Names", [BCP 32](#), [RFC 2606](#), June 1999.
- [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.
- [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.

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

Although on the surface it might appear that this document is significantly different from [[RFC3406](#)], in general it only modifies the order of presentation, with the intent of making it easier for people to define and register URN namespaces. However, the only major substantive change is removing the category of experimental namespaces, in accordance with [[RFC6648](#)].

Authors' Addresses

Peter Saint-Andre (editor)
Cisco Systems, Inc.
1899 Wynkoop Street, Suite 600
Denver, CO 80202
USA

Phone: +1-303-308-3282
Email: psaintan@cisco.com

Leslie Daigle

Dirk-Willem van Gulik

Renato Iannella

Patrick Faltstrom

