URNBIS   P. Saint-Andre, Ed.
Internet-Draft   Cisco Systems, Inc.
Obsoletes: 3406 (if approved)   L. Daigle
Intended status: BCP   Thinking Cat Enterprises
Expires: January 13, 2014   D. van Gulik
                                                           WebWeaving
                                                         R. Iannella
                                                    Semantic Identity
                                                        P. Faltstrom
                                                              Netnod
                                                      July 12, 2013

        Uniform Resource Name (URN) Namespace Definition Mechanisms
             draft-ietf-urnbis-rfc3406bis-urn-ns-reg-06

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 January 13, 2014.

Copyright Notice

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

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 the
following:

o The concept of a URN namespace.
o A mechanism for defining URN namespaces and associating each
   namespace with a public identifier (called a Namespace ID or
   "NID").
o Procedures for registering namespace identifiers 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, this document specifies
URN namespaces with the smallest reasonable set of changes from
[RFC3406]. 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?**

   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 (e.g., in the form of an algorithm).

   The common definition constraint means that both 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 some numbers might be valid
   identifiers in two different established 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, is a process that is inherently dependent on the
   requirements of the community defining the identifiers, 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.

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
   necessary 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.

4.1. 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 effectively forced someone to 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 might actively use the names
   assigned within that NID might 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 might 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 might 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 7), 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.

4.2. 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 6.

5. 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.ietf-urnbis-rfc2141bis-urn].
o The process for assigning URNs within the namespace.
o Optionally, the process for resolving URNs issued within the
   namepace.

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.

5.1. Formal Namespaces

Formal NIDs are assigned as a result of IETF Review as defined in the
"IANA Considerations" document [RFC5226]. Thus an application for a
formal NID is made by publishing an RFC in the IETF stream, either as
the product of an IETF working group or as an individual submission
sponsored by an Area Director. 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 approval
pursuant to the guidelines provided here (as well as standard RFC
publication guidelines).

5.1.1. Syntax

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

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 two characters 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.

5.1.2. Specification

The specification defining a formal namespace MUST include a
completed namespace definition template (see Section 7).

The specification also MUST include the following sections.

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

o The type of resources to be identified
o The type of services to be supported

   o Procedures for assigning URNs within this namespace

   o Processes for resolving URNs assigned within this namespace, if
      any

   It is expected that more than one namespace might 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 consideration by the
   Internet community, expert reviewers, and the IESG.

   Second, the "Community Considerations" section explains how the
   intended community will benefit by assignment 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 Methods and benefits for using the assigned URNs

   o Methods and benefits for resolving the assigned URNs (if any)

   o 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)

   Third, the "Security Considerations" section describes any potential
   security-related issues 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 lexical equivalence (see also
   [RFC6943]), leakage of private information when identifiers are
   communicated on the public Internet, the potential for directory
   harvesting, and the issues discussed in [RFC3552].

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

5.2. Informal Namespaces

   Informal namespaces are directly requested of IANA and are assigned
   based on a policy of First Come First Served [RFC5226].

   The namespace identifier assigned by IANA has the following syntax:

       "urn-" <number>

   The <number> is chosen by IANA. 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].

6. Registering a URN Namespace

6.1. Formal Namespaces

The registration policy for formal namespaces is IETF Review
[RFC5226]. The key steps for registration of a formal namespace are:

1. Submit an Internet-Draft that includes all of the information
    described under Section 5.1.2 and Section 7 of this document.
2. Send the completed namespace definition template, along with a
    pointer to the Internet-Draft, to the urn-nid@ietf.org discussion
    list for technical review.
3. If necessary to address comments received, repeat steps 1 and 2.
4. Ask the responsible Area Director to process the Internet-Draft
    for publication as an RFC. Note that the IESG can request
    further changes or direct discussion to designated working
    groups, area experts, etc.
5. If the IESG approves the document for publication as an RFC, the
    IANA will 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.

6.2. 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 7).
    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. If necessary to address comments received, repeat steps 1 and 2.
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.

7. URN Namespace Definition Template

   Definition of a URN namespace is accomplished by completing the
   following template. In addition to 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)

   Providing a complete and accurate template is particularly helpful to
   communities that are evaluating the possibility of using a portion of
   an existing URN namespace rather than creating a new namespace.

   As described under Section 5.1.2, applications for formal URN
   namespaces MUST also document the "Namespace Considerations",
   "Community Considerations", "Security 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:

         The version and date of the registration:

         - 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
               Contact information
                 (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 can be used to introduce terminology used in
          other sections. This structure can 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. If the namespace allows use of the
          URI query component, URI fragment identifier component,
          or both, such usage needs to be described here (in
          addition to any other namespace-specific syntax, such
          as distinguishers for integral parts of resources
          identified by URNs within the namespace).

          At a high level, answers might include, but are not limited to:

          - A formal definition of the structure, e.g., in terms
             of Augmented BNF for Syntax Specifications (ABNF) as
             specified in [RFC5234]
          - A regular expression for parsing the identifier into
             components, including naming authorities
          - An algorithm for generating conformant URNs
          - An explanation that the structure is opaque

     Relevant ancillary documentation:

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

          At a high level, answers might include, but are not limited to:

          - Pointers to specifications that define the syntax and
             semantics of the namespace
          - Mention of documentation that describes the processes
             followed by an organization that assigns URNs in the
             namespace
          - Explanatory material describing 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.)

At a high level, answers might 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
- Description of a method for assignment of identifiers (e.g.,
  identifiers are assigned sequentially)
- An explanation that this information is withheld (i.e.,
  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.

At a high level, answers could 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 or how to get an identifer assigned by
existing assignment authorities.

At a high level, 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).

   At a high level, answers might 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. Such rules
   ought to always have the effect of eliminating false negatives
   that might otherwise result from comparison.

   If it is appropriate and helpful to do so, reference can be
   made to the equivalence rules defined in the URI specification
   [RFC3986].

   Some examples include:

   - Equivalence between uppercase and lowercase characters in
      the Namespace Specific String
   - 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 related to handling of equivalences between
   characters in general; they are statements limited in scope to
   reflecting the rules for this specific namespace only.

Conformance with URN syntax:

This section ought to outline 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 flag 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].

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 might not be the case that all 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 might 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
resolvable using whatever mechanisms are described as part of
the namespace registration.

Validation mechanisms might be:

- A syntax grammar
- An online service
- An offline 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".

8. 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 misinformation. Information
in these declarations ought to be taken as advisory.

The definition of a URN namespace needs to account for potential
security issues related to assignment, use, and resolution of
identifiers within the namespace; see Section 5.1.2 for further
discussion.

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. References

10.1. Normative References

[I-D.ietf-urnbis-rfc2141bis-urn]
          Saint-Andre, P. and R. Moats, "Uniform Resource Name (URN)
          Syntax", draft-ietf-urnbis-rfc2141bis-urn-05 (work in
          progress), July 2013.

[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.

10.2. Informative References

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

[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.

[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.

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
interested parties to define and register URN namespaces. In
addition, some of the text was 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]. The only major substantive
change was removing the category of experimental namespaces,
consistent 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
    Thinking Cat Enterprises

    Dirk-Willem van Gulik
    WebWeaving

    Renato Iannella
    Semantic Identity

    Patrick Faltstrom
    Netnod