

Network Working Group
Request for Comments: 3406
BCP: 66
Obsoletes: 2611
Category: Best Current Practice

L. Daigle
Thinking Cat Enterprises
D.W. van Gulik
WebWeaving
R. Iannella
IPR Systems
P. Faltstrom
Cisco
October 2002

Uniform Resource Names (URN) Namespace Definition Mechanisms

Status of this Memo

This document specifies an Internet Best Current Practices for the Internet Community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

Copyright Notice

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

Abstract

This document lays out general definitions of and mechanisms for establishing Uniform Resource Names (URN) "namespaces". The URN WG has defined a syntax for URNs in RFC 2141, as well as some proposed mechanisms for their resolution and use in Internet applications in RFC 3401 and RFC 3405. The whole rests on the concept of individual "namespaces" within the URN structure. Apart from proof-of-concept namespaces, the use of existing identifiers in URNs has been discussed in RFC 2288.

Table of Contents

1.0 Introduction	2
2.0 What is a URN Namespace?	3
3.0 URN Namespace (Registration) Types	3
3.1 Experimental Namespaces	4
3.2 Informal Namespaces	4
3.3 Formal Namespaces	4
4.0 URN Namespace Registration, Update, and NID Assignment Process	6
4.1 Experimental	6
4.2 Informal	6
4.3 Formal	7
5.0 Security Considerations	9

6.0 IANA Considerations	9
7.0 References	9
Appendix A -- URN Namespace Definition Template	11
Appendix B -- Illustration	15
B.1 Example Template	15
B.2 Registration steps in practice	17
Appendix C -- Changes from RFC 2611	18
C.1 Detailed Document Changes	19
Authors' Addresses	21
Full Copyright Statement	22

1.0 Introduction

Uniform Resource Names (URNs) are resource identifiers with the specific requirements for enabling location independent identification of a resource, as well as longevity of reference. URNs are part of the larger Uniform Resource Identifier (URI) family [RFC3305] with the specific goal of providing persistent naming of resources.

There are 2 assumptions that are key to this document:

Assumption #1:

Assignment of a URN is a managed process.

I.e., not all strings that conform to URN syntax are necessarily valid URNs. A URN is assigned according to the rules of a particular namespace (in terms of syntax, semantics, and process).

Assumption #2:

The space of URN namespaces is managed.

I.e., not all syntactically correct URN namespaces (per the URN syntax definition) are valid URN namespaces. A URN namespace must have a recognized definition in order to be valid.

The purpose of this document is to outline a mechanism and provide a template for explicit namespace definition, as well as provide the mechanism for associating an identifier (called a "Namespace ID", or NID) which is registered with the Internet Assigned Numbers Authority (IANA).

Note that this document restricts itself to the description of processes for the creation of URN namespaces. If "resolution" of any so-created URN identifiers is desired, a separate process of registration in a global NID directory, such as that provided by the

DDDS system [RFC3401], is necessary. See [RFC3405] for information on obtaining registration in the DDDS global NID directory.

2.0 What is a URN Namespace?

For the purposes of URNs, a "namespace" is a collection of uniquely-assigned identifiers. That is, the identifiers are not ever assigned to more than 1 resource, nor are they ever re-assigned to a different resource. A single resource, however, may have more than one URN assigned to it for different purposes. A URN namespace itself has an identifier in order to:

- ensure global uniqueness of URNs
- (where desired) provide a cue for the structure of the identifier

For example, many identifier systems may 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. Using different designators for the two collections ensures that no two URNs will be the same for different resources (since each collection is required to uniquely assign each identifier).

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., publishing community, association of booksellers, protocol developers, etc); they are beyond the scope of the IETF URN work.

This document outlines the processes by which a collection of identifiers satisfying certain constraints (uniqueness of assignment, etc) can become a bona fide URN namespace by obtaining a NID. In a nutshell, a template for the definition of the namespace is completed for deposit with IANA, and a NID is assigned. The details of the process and possibilities for NID strings are outlined below.

3.0 URN Namespace (Registration) Types

There are three categories of URN namespaces defined here, distinguished by expected level of service and required procedures for registration. Registration processes for each of these namespace types are given in Section 4.0.

3.1 Experimental Namespaces

These are not explicitly registered with IANA. They take the form:

X-<NID>

No provision is made for avoiding collision of experimental NIDs; they are intended for use within internal or limited experimental contexts.

3.2 Informal Namespaces

These are fully fledged URN namespaces, with all the rights and requirements associated thereto. Informal namespaces can be registered in global registration services. They are required to uphold the general principles of a well-managed URN namespace -- providing persistent identification of resources, and unique assignment of identifier strings. Informal and formal namespaces (described below) differ in the NID assignment. IANA will assign an alphanumeric NID to registered informal namespaces, per the process outlined in Section 4.0.

3.3 Formal Namespaces

A formal namespace may 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, must be functional on and with the global Internet, not limited to users in communities or networks not connected to the Internet. For example, a NID that is meant for naming of physics research is requested. 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 may make use of a fee-based, privately managed, or proprietary registry for assignment of URNs in the namespace, but it may still provide benefit to some Internet users if the services associated have openly-published access protocols.

In addition to the basic registration information defined in the registration template (in Appendix A), a formal namespace request must 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, some consideration as to the longevity and maintainability of the namespace must be given. The URN WG discussed at length the issue of finding objective measures for predicting (a priori) the continued success of a namespace. No conclusion was reached -- much depends on factors that are completely beyond the technical scope of the namespace. However, the collective experience of the IETF community does contain a wealth of information on technical factors that will prevent longevity of identification. The IESG may elect not to publish a proposed namespace RFC if the IETF community consensus is that it contains technical flaws that will prevent (or seriously impair the possibility of) persistent identification.

The kinds of things the URN WG discussed included:

- the organization maintaining the URN namespace should demonstrate stability and the ability to maintain the URN namespace for a long time, and/or it should be clear how the namespace can continue to be usable/useful if the organization ceases to be able to foster it;
- it should demonstrate ability and competency in name assignment. This should improve the likelihood of persistence (e.g. to minimize the likelihood of conflicts);
- it should commit to not re-assigning existing names and 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. This does not mean that there must be resolution of such names, but that they must not resolve the name to false or stale information, and that they must not be reassigned.

These aspects, though hard to quantify objectively, should be considered by organizations/people considering the development of a Formal URN namespace, and they will be kept in mind when evaluating the technical merits of any proposed Formal namespace.

4.0 URN Namespace Registration, Update, and NID Assignment Process

Different levels of disclosure are expected/defined for namespaces. According to the level of open-forum discussion surrounding the disclosure, a URN namespace may be assigned or may request a particular identifier. The "IANA Considerations" document [RFC2434] suggests the need to specify update mechanisms for registrations -- who is given the authority to do so, from time to time, and what are the processes. Since URNs are meant to be persistently useful, few (if any) changes should be made to the structural interpretation of URN strings (e.g., adding or removing rules for lexical equivalence that might affect the interpretation of URN IDs already assigned). However, it may be important to introduce clarifications, expand the list of authorized URN assigners, etc, over the natural course of a namespace's lifetime. Specific processes are outlined below.

The official list of registered URN namespaces is maintained by IANA. URN namespace registrations are currently being posted in the anonymous FTP directory:

<http://www.iana.org/assignments/urn-namespaces>

See [RFC3232] for the current location of IANA registry.

The registration and maintenance procedures vary slightly from one namespace type (as defined in Section 3.0) to another.

4.1 Experimental

These are not explicitly registered with IANA. They take the form:

X-<NID>

No provision is made for avoiding collision of experimental NIDs; they are intended for use within internal or limited experimental contexts.

As there is no registration, no registration maintenance procedures are needed.

4.2 Informal

These are registered with IANA and are assigned a number sequence as an identifier, in the format:

"urn-" <number>

where <number> is chosen by the IANA on a First Come First Served basis (see [RFC2434]).

Registrants should send a copy of the registration template (see Appendix A), duly completed, to:

urn-nid@apps.ietf.org

and allow for a 2 week discussion period for clarifying the expression of the registration information and suggestions for technical improvements to the namespace proposal.

After suggestions for clarification of the registration information have been incorporated, the template may be submitted for assignment of a NID to:

iana@iana.org

The only restrictions on <number> are that it consist strictly of digits and that it not cause the NID to exceed length limitations outlined in the URN syntax ([RFC2141]).

Registrations may be updated by the original registrant, or an entity designated by the registrant, by updating the registration template, submitting it to the discussion list for a further 2 week discussion period, and finally resubmitting it to IANA, as described above.

4.3 Formal

Formal NIDs are assigned via IETF Consensus, as defined in [RFC2434]:

"IETF Consensus - New values are assigned through the IETF consensus process. Specifically, new assignments are made via RFCs approved by the IESG. Typically, the IESG will seek input on prospective assignments from appropriate persons (e.g., a relevant Working Group if one exists)."

Thus, the Formal NID application is made via publication of an RFC through standard IETF processes. The RFC need not be standards-track, but it will be subject to IESG review and acceptance pursuant to the guidelines written here (as well as standard RFC publication guidelines). The template defined in Appendix A may be included as part of an RFC defining some other aspect of the namespace, or it may be put forward as an RFC in its own right. The proposed template should be sent to the:

urn-nid@apps.ietf.org

mailing list to allow for a two week discussion period for clarifying the expression of the registration information, before the IESG reviews the document.

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

Considerations might include:

- URN assignment procedures
- URN resolution/delegation
- type of resources to be identified
- type of services to be supported

NOTE: 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 RFC must also include a "Community Considerations" section, which indicates the dimensions upon which the proposer expects its community to be able to 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:

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

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

A particular NID string is requested, and is assigned by IETF consensus (as defined in [RFC2434]), with the additional constraints that the NID string must:

- not be an already-registered NID
- not start with "x-" (see Type I above)
- not start with "urn-" (see Type II above)
- not start with "XY-", where XY is any combination of 2 ASCII letters (see NOTE, below)
- be more than 2 letters long

NOTE: ALL two-letter combinations, and 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 namespaces is beyond the scope of this document.

Registrations may be revised by updating the RFC through standard IETF RFC update processes (see [RFC2606] for a discussion of IETF process). In any case, a revised document, in the form of a new Internet-Draft, must be published, and the proposed updated template must be circulated on the urn-nid discussion list, allowing for a 2 week review period before pursuing publication of the new RFC document.

5.0 Security Considerations

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

6.0 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 should 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. This document defines processes to replace those outlined in [RFC2611].

7.0 References

- [ISO8601] ISO 8601 : 1988 (E), "Data elements and interchange formats - Information interchange - Representation of dates and times"
- [RFC1737] Sollins, K. and L. Masinter, "Functional Requirements for Uniform Resource Names", RFC 1737, December 1994.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
- [RFC2141] Moats, R., "URN Syntax", RFC 2141, May 1997.

- [RFC2276] Sollins, K., "Architectural Principles of Uniform Resource Name Resolution", RFC 2276, January 1998.
- [RFC2288] Lynch, C., Preston, C. and R. Daniel, "Using Existing Bibliographic Identifiers as Uniform Resource Names", RFC 2288, February 1998.
- [RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.
- [RFC2611] Daigle, L., van Gulik, D., Iannella, R. and P. Faltstrom, "URN Namespace Definition Mechanisms", RFC 2611, June 1999.
- [RFC3232] Reynolds, J, Editor, "Assigned Numbers: RFC 1700 is Replaced by an On-line Database", RFC 3232, January 2002.
- [RFC3305] Mealling, M. (Ed.) and R. Denenberg (Ed.), "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.
- [RFC3401] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part One: The Comprehensive DDDS", RFC 3401, October 2002.
- [RFC3405] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part Five: URI.ARPA Assignment Procedures", RFC 3405, October 2002.

Appendix A -- URN Namespace Definition Template

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

- entities seeking to have a URN assigned in a namespace (if applicable)
- 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.

Applications for Formal URN namespaces must also document "Namespace Considerations", "Community Considerations" and "IANA Considerations", as described in Section 4.3.

Information in the template is as follows:

Namespace ID:

Assigned by IANA. In the case of a Formal NID registration, a particular NID string may be requested.

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 outlined in [ISO8601]:

YYYY-MM-DD

Declared registrant of the namespace:

This includes:

Registering organization

Name

Address

Designated contact person

Name

Coordinates (at least one of: e-mail, phone, postal address)

Declaration of syntactic structure:

This section should 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 should be provided. If there are any specific character encoding rules (e.g., which character should always be used for single-quotes), these should 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 should 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 should address the requirement that URN identifiers be assigned uniquely -- 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 should 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 should detail the mechanisms and/or authorities for assigning URNs to resources. It should 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 must be registered in an RDS (Resolution Discovery System, see [RFC2276]) such as DDDS. Resolution then proceeds according to standard URI resolution processes, and the mechanisms of the RDS. What this section should 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; this is not relevant
- 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 should 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 should 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 [RFC2288].

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 2 issues here: 1) users should 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 telephone number-based URN namespace was created, 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 should 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" should 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".

Appendix B -- Illustration

B.1 Example Template

The following example is provided for the purposes of illustrating the URN NID template described in Appendix A. Although it is based on a hypothetical "generic Internet namespace" that has been discussed informally within the URN WG, there are still technical and infrastructural issues that would have to be resolved before such a namespace could be properly and completely described.

Namespace ID:

To be assigned

Registration Information:

Version 1

Date: <when submitted>

Declared registrant of the namespace:

Name: Thinking Cat Enterprises
Address: 1 ThinkingCat Way
Trupville, NewCountry
Contact: L. Daigle
E-mail: leslie@thinkingcat.com

Declaration of structure:

The identifier structure is as follows:

URN:<assigned number>:<FQDN>:<assigned string>

where FQDN is a fully-qualified domain name, and the assigned string is conformant to URN syntax requirements.

Relevant ancillary documentation:

Definition of domain names, found in:

P. Mockapetris, "DOMAIN NAMES - IMPLEMENTATION AND SPECIFICATION", RFC 1035, November 1987.

Identifier uniqueness considerations:

Uniqueness is guaranteed as long as the assigned string is never reassigned for a given FQDN, and that the FQDN is never reassigned.

N.B.: operationally, there is nothing that prevents a domain name from being reassigned; indeed, it is not an uncommon occurrence. This is one of the reasons that this example makes a poor URN namespace in practice, and is therefore not seriously being proposed as it stands.

Identifier persistence considerations:

Persistence of identifiers is dependent upon suitable delegation of resolution at the level of "FQDN"s, and persistence of FQDN assignment.

Same note as above.

Process of identifier assignment:

Assignment of these URNs is delegated to individual domain name holders (for FQDNs). The holder of the FQDN registration is required to maintain an entry (or delegate it) in the DDDS. Within each of these delegated name partitions, the string may be assigned per local requirements.

e.g., urn:<assigned number>:thinkingcat.com:001203

Process for identifier resolution:

Domain name holders are responsible for operating or delegating resolution servers for the FQDN in which they have assigned URNs.

Rules for Lexical Equivalence:

FQDNs are case-insensitive. Thus, the portion of the URN

urn:<assigned number>:<FQDN>:

is case-insensitive for matches. The remainder of the identifier must be considered case-sensitive.

Conformance with URN Syntax:

No special considerations.

Validation mechanism:

None specified.

Scope:

Global.

B.2 Registration steps in practice

The key steps for registration of informal or formal namespaces typically play out as follows:

Informal NID:

1. Complete the registration template. This may be done as part of an Internet-Draft.

2. Communicate the registration template to urn-nid@apps.ietf.org for technical review -- as a published I-D, or text e-mail message containing the template.
3. Update the registration template as necessary from comments, and repeat steps 2 and 3 as necessary.
4. Once comments have been addressed (and the review period has expired), send a request to IANA with the revised registration template.

Formal NID:

1. Write an Internet-Draft describing the namespace and include the registration template, duly completed. Be sure to include "Namespace Considerations", "Community Considerations" and "IANA Considerations" sections, as described in Section 4.3.
2. Send the Internet-Draft to the I-D editor, and send a copy to urn-nid@apps.ietf.org for technical review.
3. Update the Internet-Draft as necessary from comments, and repeat steps 2 and 3 as needed.
4. 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.
5. If the IESG approves the document for publication as an RFC, send a request to IANA to register the requested NID.

Appendix C -- Changes from RFC 2611

This revision of [RFC2611] adds more detail describing the process of registering a URN namespace identifier (in terms of mechanical steps).

This version of the document also separates the process (mechanics) from the discussion of the requirements for namespaces, attempting to make the latter as objective as possible.

Throughout the document, references have been updated to the current versions of the DDDS and related documentation (which collectively obsolete [RFC2168] and related drafts).

C.1 Detailed Document Changes

Added table of contents

Section 2

Clarified the definition of a URN namespace, the uniqueness of assignment, and that a single resource may have more than one identifier associated with it.

Clarified the "number example" -- that the same string may appear in 2 different namespaces, and be applied to different resources. Originally used ISBN/ISSN example, but structurally this is not possible.

Section 3 (new)

This section explicitly defines the 3 categories of namespace -- Experimental, Informal and Formal. This section provides a description of the intended use of the different namespace types, as well as some acceptability guidelines for Formal namespaces (which require IETF review).

Section 4.0

Spelled out the name of RFC 2434 ("IANA Considerations").

Provided a pointer to the IANA URN namespace registry.

Sections 4.1-4.3

New subsection divisions of the existing discussion of individual namespace types.

Section 4.2

Corrected reference to URN Syntax document (RFC 2141, not RFC 2168).

Section 4.3

Added clarifying text as to the intended nature of Formal namespaces and processes for registering them.

Added text to describe the requirement for a "Namespace Considerations" section in RFCs defining Formal namespaces. Defined the required content of that section.

Added text to describe the new requirement for a "Community Considerations" section in RFCs defining Formal namespaces. Defined the required content of that section.

Added text to explicitly call out the need for an "IANA Considerations" section in such RFCs, in order to alert IANA to required action.

Added text to further clarify the (IETF) process for revising Formal namespace registrations through the RFC and IETF review process.

Section 6

New section -- added text to describe the IANA considerations for this document.

Section 7 -- References

Added references to revised NAPTR documentation ([RFC3401]), and the previous version of this document ([RFC2611]).

Appendix A

Section created by moving the "URN Namespace Definition Template" (RFC2611's Section 3) to an appendix.

Added references to the new requirements for "Namespace Considerations", "Community Considerations", and "IANA Considerations" sections for Formal namespace registrations.

Clarified the "Declared registrant of the namespace" template element.

Added text to describe the purpose and scope of the "Validating Mechanism".

Appendix B

Section B.1 is the "example template" that was "Section 5" in RFC 2611.

Update the sample "declared registrant" data per the changes to the template description.

Removed the reference to "US-ASCII" in the "namespace specific string" of the example namespace.

Section B.2 (new)

This added section is a step-by-step walkthrough of the process for registering Informal namespaces and Formal namespaces.

Authors' Addresses

Leslie L. Daigle
Thinking Cat Enterprises
EMail: leslie@thinkingcat.com

Dirk-Willem van Gulik
WebWeaving Internet Engineering
Nieuwsteeg 37A
2311 RZ Leiden
The Netherlands

URL: <http://www.webweaving.org/>
Email: dirkx@webweaving.org

Renato Iannella
IPR Systems Pty Ltd.
EMail: renato@iprsystems.com

Patrik Faltstrom
Cisco Systems Inc
170 W Tasman Drive SJ-13/2
San Jose CA 95134
USA

EMail: paf@cisco.com

Full Copyright Statement

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

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

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

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

Acknowledgement

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