

Internet Engineering Task Force
[draft-daigle-urnframework-00.txt](#)
INTERNET-DRAFT

Leslie L. Daigle
Bunyip Information Systems Inc
Patrik Faltstrom
Tele2/Swipnet
Renato Iannella
DSTC Pty Ltd

June 13, 1996.

A Framework for the Assignment and Resolution of Uniform Resource Names

----- Abstract -----

The main purpose of URNs is to serve as persistent resource identifiers, where each URN may potentially need to be usable for several decades. A primary goal of the URN framework is to provide a mechanism whereby a URN can be resolvable for many decades, while still allowing the resolution services for that URN to be provided by the parties responsible for maintenance of the information about that URN. The Uniform Resource Name framework abstract description is proposed as a means of handling URN identifiers in modular layers of information management. From the standpoint of the bearer of a URN, this means that much of the underlying support for a given URN or collection may change without disrupting the identification capacities of the URN.

----- Status of this draft -----

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months. Internet-Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet-Drafts as reference material or to cite them other than as a ``working draft'' or ``work in progress.''

To learn the current status of any Internet-Draft, please check the `1id-abstracts.txt` listing contained in the Internet-Drafts Shadow Directories on `ds.internic.net`, `nic.nordu.net`, `ftp.isi.edu`, or `munniari.oz.au`.

This draft expires November 13, 1996.

Updated versions and related information can be found at:

<http://www.bunyip.com/research/ietf/urn-bof>

Introduction

The main purpose of URNs is to serve as persistent resource identifiers, where each URN may potentially need to be usable for several decades. A primary goal of the URN framework is to provide a mechanism whereby a URN can be resolvable for many decades, while still allowing the resolution services for that URN to be provided by the parties responsible for maintenance of the information about that URN. Since the resources named by a URN may change ownership over time, the resolution services for a URN must be able to change without changes to the URN itself.

The Uniform Resource Naming framework defines the structures and services required in order to support uniform resource naming across a number of different namespaces. The underlying characteristics of these namespaces may vary considerably, but the URN framework lays out the basic requirements a namespace must (claim to) live up to in order to be supported as a partition of the URN namespace. When this work is complete, these requirements will exist in the form of an objectively verifiable checklist against which new namespaces can be compared/supporting structures built.

The URN framework aims to provide flexible methods to support various 'resolution services' for URNs. The resolution service will provide a client with access to the named resource. The URN framework will provide the client with the ability to instantiate these resolution services, in a manner that will not bind any client or service to protocol dependant technologies.

The URN framework abstract description is proposed as a means of handling the URN identifiers in modular layers of information management. From the standpoint of the bearer of a URN, this means that much of the underlying support for a given URN or collection may change without disrupting the identification capacities of the URN. For example, the resolution service(s) used by a collection may move or change format entirely, without affecting the validity of the URN.

Beyond accommodating differences in individual namespaces, this framework is also designed to modularize authority in such a way as to keep information maintenance close to the responsible entity. For example, if a resource location changes, only the final level of resolution needs to be made aware of the change.

Syntax

For the purposes of this document, the following informal syntax is used:

URN:<NID>:<NSS>

NID is the Namespace Identifier and NSS is the Namespace Specific String. Each namespace defines the structure of its NSS, and the Namespace ID is used to determine the `_syntactic_` interpretation of the Namespace Specific String to the extent of extracting the Naming Authority information. This may be done by providing rules for extracting specific Naming Authority ID information from the NSS, or by having a specified set of Naming Authorities for an entire namespace. The NID[+NA] serves as a key to finding the location of resolution services. The protocol used to resolve the name is independent of the name itself. This construct allows the incorporation of different name structures.

Examples:

URN:duns:002372413:annual-report-1997

In this example, "duns" is the NID, and "002372413:annual-report-1997" is the NSS.

URN:inet:gatech.edu:olympic-schedule-1997

Here, "inet" is the NID, and "gatech.edu:olympic-schedule-1997" is the NSS.

Each of the "duns" and "inet" namespaces define rules for the syntactic interpretation of the NSS to the extent of extracting an NA string.

URN Resolution through Registries

The URN framework provides URN resolution infrastructure in the shape of an abstract description of registries of namespaces and naming authorities. Resolution of a URN, when possible, can be accomplished by consulting these registries, although client software may implement strategies for optimization of the process.

This default method of resolution-with-registries proceeds as follows:

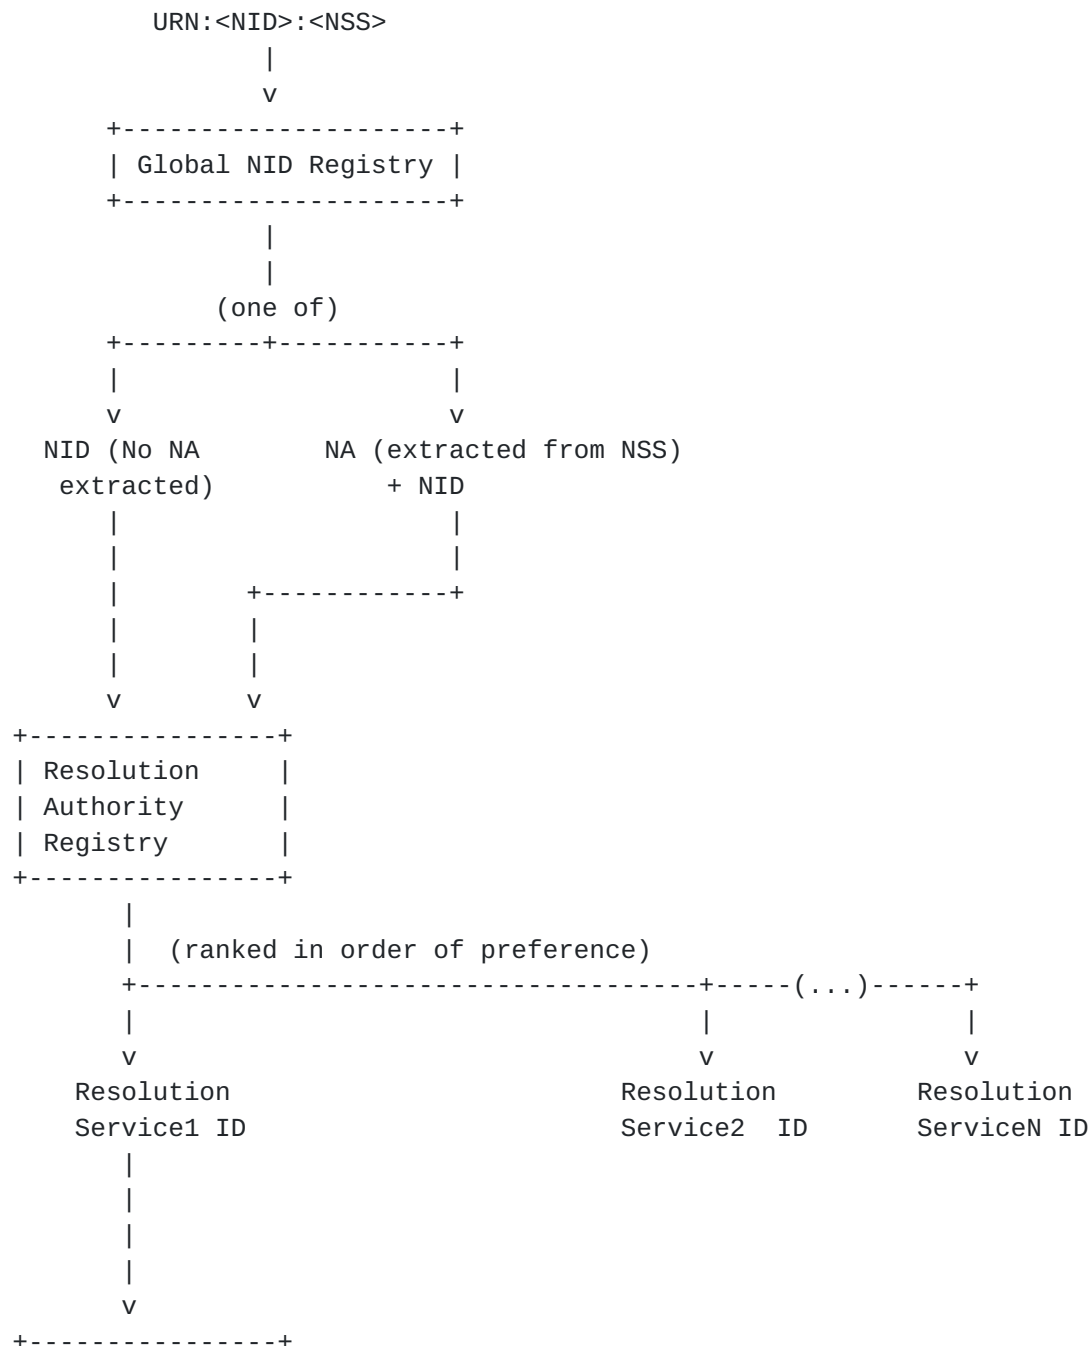
- . A NID registry consulted get a description of the NSS. (See [draft-daniel-naptr-01.txt](#) for a description of the implementation of such an NID registry).
- . If there is a NA string in the NSS, the next step of the resolution is to look up the NID+NA; otherwise the NID is used again in this step of resolution. The lookup occurs in a resolution


authority registry (see [draft-daniel-naptr-01.txt](#) for a description of the implementation of such a registry) to identify one or more resolution services for the result.

- . The URN is passed to one of the resolution services.

Note that this document refers to "a global NID registry". There is still discussion underway as to the best means of handling public and private namespaces and different ways of implementing registries.

Pictorially, this can be represented as:





N.B.: <result> is defined abstractly -- it might be the desired resource, a description of the resource, the location of the resource, etc.

NID registry:

While there needs to be some form of global NID registry (which is not an uncommon requirement for Internet infrastructure, e.g., the DNS service), it is possible that a given client will refer first (or exclusively) to a local NID registry. This will certainly be useful for closed software information systems.

It is expected that the NID registry will be fairly small, and pretty involatile. It is expected that caching will be quite successful in handling the majority of requests.

Some namespaces will not want to publish rules for parsing names -- i.e., the (partial) decoding of the NSS might be considered violation of the proprietary nature of the name assignment scheme, giving the world at large more information about the information in the namespace than the namespace's owner is happy about. Such namespaces can either opt to pre-pend other indications of naming authorities in translating their names into NSS's, or these will fall into the category of URN's from which the NA cannot be extracted.

Resolution Authority Registry:

Once the NA (if any) has been extracted from the URN, the client consults the Resolution Authority Registry to determine who can resolve names issued in that namespace by that naming authority. Depending on the client software, there may be more than one resolution authority registry that can be consulted.

Examples:

URN:duns:002372413:annual-report-1997

If the duns namespace does not provide rules for extracting the NA information from the NSS, NA's for the whole namespace must be provided.

URN:inet:gatech.edu:olympic-schedule-1997

The inet namespace may provide rules for extracting "gatech.edu" as the key for composing an NA that should be looked up in the Naming Authority Registry.

Similarly,

URN:message-id:199606071447.KAA20718@cs.utk.edu

The message-id namespace may provide rules for extracting "cs.utk.edu" as the key for composing an NA that should be looked up in the Naming Authority Registry.

Clients can do what they want...

In providing the above as the default URN resolution procedure, it is expected that there will be at least one "system" (infrastructure) deployed to support it -- some global name registry, and a set of NA identification servers. Registration in these services is required as a step in becoming an official URN-supported namespace. This is required in order to be able to ensure global fall-back mechanisms for resolution of orphaned or migrated resources.

Nevertheless, the resolution strategy used by any client is chosen by the developers of that client. For example, a given application may choose to provide and support a directory of often-used URNs, and its client software may first consult that directory. Defining the reliability and optimality of such strategies is beyond the scope of the URN framework; may the best strategy win. The global infrastructure that is provided by the URN framework is designed to ensure the overall goals of URNs -- global scope, persistence, etc.

The URN Framework
-- what is in, what is not

This proposed URN framework addresses the URN requirements as laid out in RFC [1737](#). **Some of these are provided by the infrastructure of the framework itself**, some are defined as characteristics a given namespace must support in order to be considered eligible to participate in the URN framework, and a few are considered to lie within the realm of individual namespaces (e.g., characteristics that are important to one community, but that impede the implementation of characteristics necessary to another).

In particular, the URN framework provides global scope, allows persistence, independence, and legacy support. Participating Namespaces must claim to provide global uniqueness in the URNs generated in order to be part of the

URN framework, although they may vary in how they provide persistence and scalability.

Namespace Identifier Requirements and
Administration

Some set of criteria must be established to govern what can and cannot be used as a (top-level) namespace (NID assignment). This can include:

- . demonstration of an established namespace management system (including an overview of mechanisms for preventing the duplication/reassignment of name identifiers. These mechanisms are particular to the individual namespaces).
- . multi-organization participation in the naming system
- . rules on how names are assigned, name assignment authority delegation
- . provision of escape clause

Related issues

The URN framework, described here in its abstract form, is meant to provide the infrastructure for a system that can fulfill the desiderata of URN systems. This section describes certain stated issues that are handled indirectly by the URN framework -- that is, the hooks are available in the abstract framework, but the implementation will be dependent on elements that are outside the scope of the URN infrastructure per se.

Lexical Equivalence

This definition of the abstract URN framework does not immediately provide mechanisms to determine equivalence of two URNs. That is, in order to accommodate the variety of existing and proposed namespaces, it is not possible to define a single scheme that would incorporate this ability. Nevertheless, there is every likelihood that single resources may be assigned more than one URN (under different naming schemes, or as the resources are included in different perspectives of a collection, for instance). It is expected that, where needed, individual services for determining "equivalence" will evolve.

Authority Over Name Strings

The modular nature of the framework distributes authority across different levels in the resolution process. Therefore, while the Global NID

Registry will provide the official response about what Resolution Authority Registries are registered to handle a given URN, it has nothing to say about the NSS. Similarly, individual Resolution Services maintain authority for resolving the URN to a particular result, but they have made no authoritative statements about a namespace.

Rules for individual namespaces are determined by the owners of the namespace before they are registered in the NID registry. Authority of assignment of individual names within the namespace is also dependent on the individual namespace, and is not part of the URN resolution framework.

Scope of Grandfathering

Grandfathering of existing namespaces under the URN framework will be handled in the same way as the creation of new partitions -- a proposal must be put forward to demonstrate that the namespace can fulfill the requirements of URN namespace partitions (e.g., generating globally unique URN identifiers) and what rules will be registered in the Global NID Registry.

Acknowledgements

The URN work was initiated as part of the IETF URI Working Group. When that group was closed at the Stockholm IETF in July, 1995, interested parties continued to meet to move the work forward.

In addition to the URN BOF held at the Dallas IETF in December, 1995, three voluntary meetings have been organized (in Knoxville, Montreal, and Los Alamos). This document is a direct product of those discussions, and the following attendees must be thanked for their hard work in helping move this effort forward:

Ron Daniel
David Ely
Dirk van Gulik
Dan Laliberte
Michael Mealling
Keith Moore
Keith Shafer
Michael Shapiro
Reed Wade
Stuart Weibel
Chris Weider
Ted Wolf

Authors' Addresses

Leslie L. Daigle
Bunyip Information Systems Inc.
310 Ste. Catherine St. West
Suite 300
Montreal, Quebec, Canada
H2X 2A1
voice: (514) 875-8611
fax: (514) 875-8134
e-mail: leslie@bunyip.com

Patrik Faltstrom
Tele2/Swipnet
Box 62
Borgarfjordsgatan 16
S-164 94 Kista
Sweden
Phone: +46-8-56264000
Fax: +46-8-56264200
Email: paf@swip.net

Renato Iannella
Research Data Network CRC
DSTC Pty Ltd
Gehrmann Laboratories
University of Queensland
Australia
Phone: +61-7--3365-4310
Fax: +61-7--3365-4311
Email: renato@dstc.edu.au