

EXHIBIT 16


<https://www.ietf.org/rfc/rfc1737.txt>

Go

[148 captures](#)

13 Oct 1999 - 19 Dec 2017

Network Working Group
Request for Comments: 1737
Category: Informational

K. Sollins
MIT/LCS
L. Masinter
Xerox Corporation
December 1994

Functional Requirements for Uniform Resource Names

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

1. Introduction

This document specifies a minimum set of requirements for a kind of Internet resource identifier known as Uniform Resource Names (URNs). URNs fit within a larger Internet information architecture, which in turn is composed of, additionally, Uniform Resource Characteristics (URCs), and Uniform Resource Locators (URLs). URNs are used for identification, URCs for including meta-information, and URLs for locating or finding resources. It is provided as a basis for evaluating standards for URNs. The discussions of this work have occurred on the mailing list uri@bunyip.com and at the URI Working Group sessions of the IETF.

The requirements described here are not necessarily exhaustive; for example, there are several issues dealing with support for replication of resources and with security that have been discussed; however, the problems are not well enough understood at this time to include specific requirements in those areas here.

Within the general area of distributed object systems design, there are many concepts and designs that are discussed under the general topic of "naming". The URN requirements here are for a facility that addresses a different (and, in general, more stringent) set of needs than are frequently the domain of general object naming.

The requirements for Uniform Resource Names fit within the overall architecture of Uniform Resource Identification. In order to build applications in the most general case, the user must be able to discover and identify the information, objects, or what we will call in this architecture resources, on which the application is to operate. Beyond this statement, the URI architecture does not define "resource." As the network and interconnectivity grow, the ability to make use of remote, perhaps independently managed, resources will

- o It is strongly recommended that there be a mapping between the names generated by each naming authority and URLs. At any specific time there will be zero or more URLs into which a particular URN can be mapped. The naming authority itself need not provide the mapping from URN to URL.
- o For URNs to be transcribable and transported in mail, it is necessary to limit the character set usable in URNs, although there is not yet consensus on what the limit might be.

In assigning names, a name assignment authority must abide by the preceding constraints, as well as defining its own criteria for determining the necessity or indication of a new name assignment.

5. Other considerations

There are three issues about which this document has intentionally not taken a position, because it is believed that these are issues to be decided by local determination or other services within an information infrastructure. These issues are equality of resources, reflection of visible semantics in a URN, and name resolution.

One of the ways in which naming authorities, the assigners of names, may choose to make themselves distinctive is by the algorithms by which they distinguish or do not distinguish resources from each other. For example, a publisher may choose to distinguish among multiple printings of a book, in which minor spelling and typographical mistakes have been made, but a library may prefer not to make that distinction. Furthermore, no one algorithm for testing for equality is likely to be applicable to all sorts of information. For example, an algorithm based on testing the equality of two books is unlikely to be useful when testing the equality of two spreadsheets. Thus, although this document requires that any particular naming authority use one algorithm for determining whether two resources it is comparing are the same or different, each naming authority can use a different such algorithm and a naming authority may restrict the set of resources it chooses to identify in any way at all.

A naming authority will also have some algorithm for actually choosing a name within its namespace. It may have an algorithm that actually embeds in some way some knowledge about the resource. In turn, that embedding may or may not be made public, and may or may not be visible to potential clients. For example, an unreflective URN, simply provides monotonically increasing serial numbers for resources. This conveys nothing other than the identity determined by the equality testing algorithm and an ordering of name assignment by this server. It carries no information about the resource itself.

An MD5 of the resource at some point, in and of itself may be reflective of its contents, and, in fact, the naming authority may be perfectly willing to publish the fact that it is using MD5, but if the resource is mutable, it still will be the case that any potential client cannot do much with the URN other than check for equality. If, in contrast, a URN scheme has much in common with the assignment ISBN numbers, the algorithm for assigning them is public and by knowing it, given a particular ISBN number, one can learn something more about the resource in question. This full range of possibilities is allowed according to this requirements document, although it is intended that naming authorities be discouraged from making accessible to clients semantic information about the resource, on the assumption that that may change with time and therefore it is unwise to encourage people in any way to depend on that semantics being valid.

Last, this document intentionally does not address the problem of name resolution, other than to recommend that for each naming authority a name translation mechanism exist. Naming authorities assign names, while resolvers or location services of some sort assist or provide URN to URL mapping. There may be one or many such services for the resources named by a particular naming authority. It may also be the case that there are generic ones providing service for many resources of differing naming authorities. Some may be authoritative and others not. Some may be highly reliable or highly available or highly responsive to updates or highly focussed by other criteria such as subject matter. Of course, it is also possible that some naming authorities will also act as resolvers for the resources they have named. This document supports and encourages third party and distributed services in this area, and therefore intentionally makes no statements about requirements of URNs or naming authorities on resolvers.

Security Considerations

Applications that require translation from names to locations, and the resources themselves may require the resources to be authenticated. It seems generally that the information about the authentication of either the name or the resource to which it refers should be carried by separate information passed along with the URN rather than in the URN itself.