

SIPPING
Internet-Draft
Expires: April 26, 2006

H. Schulzrinne
Columbia U.
October 23, 2005

A Uniform Resource Name (URN) for Services
draft-schulzrinne-sipping-service-01

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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 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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on April 26, 2006.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

The content of many communication services depend on the context, such as the user's location. We describe a 'service' URN that allows to register such context-dependent services that can be resolved in a distributed manner.

Internet-Draft

Service URN

October 2005

Table of Contents

- [1.](#) Introduction [3](#)
- [2.](#) Registration Template [4](#)
- [3.](#) Example [5](#)
- [4.](#) IANA Considerations [6](#)
- [5.](#) References [6](#)
 - [5.1](#) Normative References [6](#)
 - [5.2](#) Informative References [6](#)
- Author's Address [7](#)
- [A.](#) Acknowledgments [7](#)
 - Intellectual Property and Copyright Statements [8](#)

Internet-Draft

Service URN

October 2005

1. Introduction

In existing telecommunications systems, there are many well-known communication and information services that are offered by loosely coordinated entities across a large geographic region, with well-known identifiers. Some of the services are operated by governments or regulated monopolies, others by competing commercial enterprises. Examples include emergency services (reached by 911 in North America, 112 in Europe), telephone directory and repair services (411 and 611 in the United States and Canada), government information services (311 in some cities in the United States), lawyer referral services (1-800-LAWYER), car roadside assistance (automobile clubs) and pizza delivery services. Unfortunately, almost all of them are limited in scope to a single country or possibly a group of countries, such as those belonging to the North American Numbering Plan or the European Union. The same identifiers are often used for other purposes outside that region, making accessing such services difficult when users travel or use devices produced outside their home country.

These services are characterized by long-term stability of user-visible identifiers, decentralized administration of the underlying service and a well-defined resolution mechanism. (For example, there is no national coordination or call center for 911; rather, various local government organizations cooperate to provide this service, based on jurisdictions.)

In this document, we propose a URN namespace that, together with resolution protocols beyond the scope of this document, allows to define such global, well-known services, while distributing the actual implementation across a large number of service-providing entities. While there are many ways to divide provision of such services, we focus on geography as a common way to delineate service regions. In addition, users can choose different directory providers that in turn manage how geographic locations are mapped to service providers.

Availability of such service identifiers simplifies end system configuration. For example, an IP phone could have a special set of short cuts or buttons that invoke emergency services, as it would not be practical to manually re-configure the device with local emergency contacts for each city or town a user visits with his or her mobile device. Also, such identifiers allow to delegate routing decisions to third parties and mark certain requests as having special characteristics while preventing these characteristics to be accidentally invoked on inappropriate requests.

This URN allows to identify services independent of a particular protocol to deliver the services. It may appear in protocols that

allow general URIs, such as SIP [\[4\]](#) request URIs, web pages or mapping protocols.

Existing technologies address the mapping of service identifiers to a service for a particular DNS domain (DNS SRV [\[6\]](#), DNS NAPTR [\[7\]](#)) or a local area network (SLP [\[5\]](#)).

The tel URI [\[9\]](#) allows to express service codes such as 911 by adding a context parameter, but does not address the problem of global validity.

LUMP [\[10\]](#) is a prototype resolution system for mapping URNs to URLs based on geographic location. However, it is anticipated that there will be several such systems.

[2.](#) Registration Template

Below, we include the registration template for the URN scheme according to [RFC 3406](#) [\[8\]](#).

Namespace ID: service

Registration Information: Registration version: 1; registration date: 2005-07-10

Declared registrant of the namespace: TBD

Declaration of syntactic structure: The URN consists of a hierarchical service identifier, with a sequence of labels separated by periods. The left-most label is the most significant one and is called 'top-level service', while names to the right are called 'sub-services'. The set of allowable characters is the same as that used for domain names. Any string of service labels

can be used to request services that are either more generic or more specific. In other words, if a service 'x.y.z' exists, the URNs 'x' and 'x.y' are also valid service URNs. [?]

```
"URN:service:" top-level-service *("." service-identifier)
top-level-service = ALPHA / DIGIT / "-" /
service-identifier = ALPHA / DIGIT / "-" /
```

Relevant ancillary documentation: None

Identifier uniqueness considerations: 'service' URNs identify one logical service, recognized by human users as such. The service does not have to be provided by the same organization or to the same standards over time and space. Unlike for other URNs, the content of the service is by nature dynamic. While undesirable in many cases, two users making the same request for a service from the same place may not necessarily be directed to the same resource.

Identifier persistence considerations: The 'service' URN for the same service is expected to be persistent, although there naturally cannot be a guarantee that a particular service will continue to be available globally or at all times.

Process of identifier assignment: Details of the service assignment depend on the service and national regulations. In general, it is assumed that providers of services can register through a service mapping mechanism for a particular service in a particular geographic area. The provision of some services may be restricted by local or national regulations. (As a hypothetical example, providing emergency services may be restricted to government-authorized entities, which may limit the region where each entity can advertise its services.) The rules for each service are described in a service-specific document.

Process for identifier resolution: 'service' identifiers are resolved by the TBD mapping protocol, an instance of a Resolution Discovery System (RDS) as described in [RFC 2276](#) [2]. (In theory, there could be several such mapping protocols in concurrent use, as long as there are reasonable guarantees that all services are available in all mapping protocols.)

Rules for Lexical Equivalence: 'service' identifiers are compared according to domain name comparison rules. The use of homographic

identifiers is NOT RECOMMENDED.
Conformance with URN Syntax: There are no special considerations.
Validation mechanism: The RDS mechanism is also used to validate the existence of a resource. As noted, by its design, the availability of a resource may depend on where service is desired and there may not be service available in all or most locations. (For example, roadside assistance service is unlikely to be available on about 70% of the earth's surface.)
Scope: The scope for this URN is public and global.

[3.](#) Example

For discussion and illustration purposes only, we include an example of a particular service. We choose emergency services as an example, with the top-level service identifier 'sos'. A possible list of identifiers might include:

```
urn:service:sos
urn:service:sos.fire
urn:service:sos.police
urn:service:sos.marine
urn:service:sos.mountain
urn:service:sos.rescue
urn:service:sos.poison
urn:service:sos.suicide
urn:service:sos.mental-health
```

[4.](#) IANA Considerations

New service-identifying tokens and sub-registrations are to be managed by IANA, according to the processes outlined in [\[3\]](#). The policy for top-level service names is TBD, but could be 'specification required', 'IETF Consensus' or 'Standards Action'. The policy for assigning names to sub-services may differ for each top-level service designation and MUST be defined by the document describing the top-level service.

[5.](#) References

[5.1](#) Normative References

[1] Bradner, S., "Key words for use in RFCs to Indicate Requirement

- Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [2] Sollins, K., "Architectural Principles of Uniform Resource Name Resolution", [RFC 2276](#), January 1998.
- [3] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 2434](#), October 1998.
- [4] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", [RFC 3261](#), June 2002.

[5.2](#) Informative References

- [5] Guttman, E., Perkins, C., Veizades, J., and M. Day, "Service Location Protocol, Version 2", [RFC 2608](#), June 1999.
- [6] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", [RFC 2782](#), February 2000.
- [7] Mealling, M. and R. Daniel, "The Naming Authority Pointer (NAPTR) DNS Resource Record", [RFC 2915](#), September 2000.
- [8] Daigle, L., van Gulik, D., Iannella, R., and P. Faltstrom, "Uniform Resource Names (URN) Namespace Definition Mechanisms", [BCP 66](#), [RFC 3406](#), October 2002.
- [9] Schulzrinne, H., "The tel URI for Telephone Numbers", [RFC 3966](#), December 2004.
- [10] Schulzrinne, H., "Location-to-URL Mapping Protocol (LUMP)", [draft-schulzrinne-ecrit-lump-00](#) (work in progress), June 2005.

Author's Address

Henning Schulzrinne
Columbia University
Department of Computer Science
450 Computer Science Building
New York, NY 10027
US

Phone: +1 212 939 7004
Email: hgs+simple@cs.columbia.edu
URI: <http://www.cs.columbia.edu>

Appendix A. Acknowledgments

This document is based on discussions with Jonathan Rosenberg.

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Copyright Statement

Copyright (C) The Internet Society (2005). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Acknowledgment

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