ECRIT
Internet-Draft
Expires: October 4, 2006

H. Schulzrinne
Columbia U.

April 2, 2006

A Uniform Resource Name (URN) for Services
draft-ietf-ecrit-service-urn-02

Status of this Memo

Copyright Notice

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.

Table of Contents

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 dialing 911 in North America, 112 in Europe), community services and volunteer opportunities (211 in some regions of the United States),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 "9-1-1" in the United States; 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. There are many ways to divide provision of such services, such as dividing responsibility by geographic region or by the service provider a user chooses. 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 the Session Initiation Protocol (SIP) [5] request URIs, web pages or mapping protocols.

The service URN is generally not expected to be visible to humans. For example, it is expected that callers will still dial '9-1-1' in the United States to reach emergency services. In some other cases, speed dial buttons might identify the service, as is common practice on hotel phones today. (Speed dial buttons for summoning emergency help are considered inappropriate by most emergency services professionals, at least for mobile devices, as they are too prone to being triggered accidentally.) Rather, protocol elements would carry the service URN described here, allowing universal identification. The translation of dial strings to service URNs is beyond the scope of this document; it is likely to depend on the location of the caller and may be many-to-one. For example, a phone for a traveler could recognize the emergency dial string for both the traveler's home location and the traveler's visited location, translating both to the same universal service URN, urn:service:sos.

Since service URNs are not routable, a proxy or user agent has to translate the service URN into a routable URI for a location-appropriate service provider, such as a SIP URL. LoST [20] is one resolution system for mapping service URNs to URLs based on geographic location. It is anticipated that there will be several such systems, possibly with different systems for different services.

We discuss alternative approaches, and why they are unsatisfactory, in Appendix A.

2. Registration Template

Below, we include the registration template for the URN scheme according to RFC 3406 [15].

Namespace ID: service

Registration Information: Registration version: 1; registration date: 2006-04-02

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 for domain names [1] and a subset of the labels allowed in [6]; labels are case-insensitive and SHOULD be specified in all lower-case. 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:" service
service            = top-level *("." service-identifier)
let-dig            = ALPHA / DIGIT
let-dig-hyp        = let-dig / '-'
service-identifier = let-dig [ *let-dig-hyp let-dig ]
top-level          = let-dig [ *25let-dig-hyp let-dig ]
```

Relevant ancillary documentation: None

Community considerations: The service URN is believe to be relevant to a large cross-section of Internet users, including both technical and non-technical users, on a variety of devices, but particularly for mobile and nomadic users. The service URN will allow Internet users needing services to identify the service by kind, without having to determine manually who provides the particular service in the user's current context, e.g., at his current location. For example, a traveler will be able to use his mobile device to request emergency services without having to know the local emergency number. The assignment of identifiers is described in the IANA Considerations (Section 4). The service URN does not prescribe a particular resolution mechanism, but it is assumed that a number of different entities could operate and offer such mechanisms.

Namespace considerations: There do not appear to be other URN namespaces that serve the same need of uniquely identifying widely-available communication and information services. Unlike most other currently registered URN namespaces, the service URN does not identify documents and protocol objects (e.g., [13], [14], [18], [19]), types of telecommunications equipment [17], people or organizations [12]. tel URIs [16] identify telephone numbers, but numbers commonly identifying services, such as 911 or 112, are specific to a particular region or country.

Identifier uniqueness considerations: A service URN identifies a logical service, specified in the service registration (see IANA considerations). Resolution of the URN, if successful, will return a particular instance of the service, and this instance may be different even for two users making the same request in the same place at the same time; the logical service identified by the URN, however, is persistent and unique. Service URNs MUST be unique for each unique service; this is guaranteed through the registration of each service within this namespace, described in Section 4.

Identifier persistence considerations: The 'service' URN for the same service is expected to be persisent, 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: The process of identifier assignment is described in the IANA Considerations (Section 4).

Process for identifier resolution: 'service' identifiers are resolved by the mapping protocols, an instance of a Resolution Discovery System (RDS) as described in RFC 2276 [3]. Each top-level service can provide its own distinct set of mapping protocols. Within each top-level service, all mapping protocols MUST return the same set of mappings. Section 3 describes how DNS NAPTR records are used to find an instance of a mapping service.

Rules for Lexical Equivalence: 'service' identifiers are compared according to case-insensitive string equality.

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. Finding the Mapping Server

When a network entity receives a service URN, it uses the S-NAPTR [6] mechanism to determine how to map the service URN, possibly using other information such as geographic location, to a routable URI. Each top-level service may define one or more such mapping protocols and mapping protocol servers may be operated by a range of providers. Thus, the network entity that needs to resolve the service URN queries an appropriate domain, typically its home or service provider domain, for NAPTR records and then selects records that match the service and the mapping protocols it supports. The application service for this URN is registered in IANA Considerations (Section 4) of this document; the application protocols are registered in the appropriate protocol document.

The S-NAPTR entry MAY contain the "s" flag if the resolving client needs to perform an SRV resolution on the replacement string.

The first entry in the following example indicates that 'sos' service URNs should be mapped to URIs using the LoST [20] protocol server at lost.example.com, a DNS A record. The second entry is for an imaginary top-level service 'pizza', using the equally imagined 'Pizza Location Protocol', offered by the pizzahouse.example.net server, which should be queried for the appropriate DNS SRV record.

Note that these NAPTR records are maintained by example.com, i.e., example.com does not actually provide the mapping service itself.

```
example.com.
;      order pref flags service      regexp
IN NAPTR 50   50  "a"  "SURN.sos:LoST"     ""
;  replacement
   lost.example.org

IN NAPTR 10 50 "s" "SURN.pizza:PLP"  ""
   _plp._tcp.pizzahouse.example.net
```

4. IANA Considerations

4.1 New Service-Identifying Tokens

New service-identifying tokens and sub-registrations are to be managed by IANA, according to the processes outlined in [4]. The policy for top-level service names is 'IETF Consensus'. 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.

To allow use within the constraints of S-NAPTR [6], all top-level service names MUST NOT exceed 27 characters.

4.2 S-NAPTR Application Service Label

Since each top-level service could use one or more different resolution protocols, we need to indicate the top-level service in the S-NAPTR application service label. To indicate the URN-to-service mapping service, all such services start with the string "SURN." (for "service URN"), followed by the top-level service identifier. Note that application service labels are case-insensitive and rendered here in mixed case purely for readability.

This document registers the label "SURN.sos" as the S-NAPTR application service label according to [6] for emergency services and defines the intended usage, interoperability considerations and security considerations (Section 5).

4.3 sos Service Types

The 'sos' service type describes emergency services and services related to public safety and health, typically offered by various branches of the government or other public institutions. Additional

sub-services can be added after expert review and should be of general public interest.

urn:service:sos The generic 'sos' service reaches a public safety answering point (PSAP), that in turn dispatches aid appropriate to the emergency. It encompasses all of the services listed below.
urn:service:sos.ambulance This service identifier reaches an ambulance service that provides emergency medical assistance and transportation.
urn:service:sos.animal-control Animal control is defined as control of dogs, cats, and domesticated or undomesticated animals.
urn:service:sos.fire The 'fire' service identifier summons the fire service, also known as the fire brigade or fire department.
urn:service:sos.gas The 'gas' service allows the reporting of natural gas (and other flammable gas) leaks or other natural gas emergencies.
urn:service:sos.mountain The 'mountain' service refers to mountain rescue services, i.e., search and rescue activities that occur in a mountainous environment, although the term is sometimes also used to apply to search and rescue in other wilderness environments.
urn:service:sos.marine The 'marine' service refers to maritime search and rescue services such as those offered by the coast guard, lifeboat or surf lifesavers.
urn:service:sos.physician The 'physician' emergency service connects the caller to a physician referral service.
urn:service:sos.poison The 'poison' service refers to special information centers set up to inform citizens about how to respond to potential poisoning. These poison control centers maintain a database of poisons and appropriate emergency treatment.
urn:service:sos.police The 'police' service refers to the police department or other law enforcement authorities.
urn:service:sos.suicide The 'suicide' service refers to the suicide prevention hotline.
urn:service:sos.mental-health The 'mental-health' service refers to the "Diagnostic, treatment, and preventive care that helps improve how persons with mental illness feel both physically and emotionally as well as how they interact with other persons." (U.S. Department of Health and Human Services)

5. Security Considerations

As an identifier, the service URN does not appear to raise any particular security issues. The services described by the URN are meant to be well-known, even if the particular service instance is access-controlled, so privacy considerations do not apply to the URN. There are likely no specific privacy issues when including a service URN on a web page, for example. On the other hand, ferrying the URN

in a signaling protocol can give attackers information on the kind of service desired by the caller. For example, this makes it easier for the attacker to automatically find all calls for emergency services or directory assistance. Appropriate, protocol-specific security mechanisms need to be implemented for protocols carrying service URNs. The mapping protocol needs to address a number of threats, as detailed in [21].

6. References

6.1 Normative References

[1] Braden, R., "Requirements for Internet Hosts - Application and Support", STD 3, RFC 1123, October 1989.

[2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[3] Sollins, K., "Architectural Principles of Uniform Resource Name Resolution", RFC 2276, January 1998.

[4] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

[5] 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.

[6] Daigle, L. and A. Newton, "Domain-Based Application Service Location Using SRV RRs and the Dynamic Delegation Discovery Service (DDDS)", RFC 3958, January 2005.

6.2 Informative References

[7] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[8] Crocker, D., "MAILBOX NAMES FOR COMMON SERVICES, ROLES AND FUNCTIONS", RFC 2142, May 1997.

[9] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, February 2000.

[10] Resnick, P., "Internet Message Format", RFC 2822, April 2001.

[11] Mealling, M. and R. Daniel, "The Naming Authority Pointer (NAPTR) DNS Resource Record", RFC 2915, September 2000.

[12] Mealling, M., "The Network Solutions Personal Internet Name (PIN): A URN Namespace for People and Organizations", RFC 3043, January 2001.

[13] Rozenfeld, S., "Using The ISSN (International Serial Standard Number) as URN (Uniform Resource Names) within an ISSN-URN Namespace", RFC 3044, January 2001.

[14] Hakala, J. and H. Walravens, "Using International Standard Book Numbers as Uniform Resource Names", RFC 3187, October 2001.

[15] Daigle, L., van Gulik, D., Iannella, R., and P. Faltstrom, "Uniform Resource Names (URN) Namespace Definition Mechanisms", BCP 66, RFC 3406, October 2002.

[16] Schulzrinne, H., "The tel URI for Telephone Numbers", RFC 3966, December 2004.

[17] Tesink, K. and R. Fox, "A Uniform Resource Name (URN) Namespace for the Common Language Equipment Identifier (CLEI) Code", RFC 4152, August 2005.

[18] Kang, S., "Using Universal Content Identifier (UCI) as Uniform Resource Names (URN)", RFC 4179, October 2005.

[19] Kameyama, W., "A Uniform Resource Name (URN) Namespace for the TV-Anytime Forum", RFC 4195, October 2005.

[20] Hardie, T., "LoST: A Location-to-Service Translation Protocol", draft-hardie-ecrit-lost-00 (work in progress), March 2006.

[21] Taylor, T., "Security Threats and Requirements for Emergency Call Marking and Mapping", draft-ietf-ecrit-security-threats-00 (work in progress), March 2006.

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+ecrit@cs.columbia.edu
URI: http://www.cs.columbia.edu

Appendix A. Alternative Approaches Considered

The "sos" SIP URI reserved user name proposed here follows the convention of RFC 2142 [8] and the "postmaster" convention documented in RFC 2822 [10]. The approach has the advantage that only the home proxy for a user needs to understand the convention and that the mechanism is likely backwards-compatible with most SIP user agents, with the only requirement that they have to be able to generate alphanumeric URLs. One drawback is that it may conflict with locally assigned addresses of the form "sos@domain". Also, if proxies not affiliated with the domain translate the URL, they violate the current SIP protocol conventions.

There are a number of possible alternatives, each with their own set of advantages and problems:

tel:NNN;context=+C This approach uses tel URIs [16]. Here, NNN is the national emergency number, where the country is identified by the context C. This approach is easy for user agents to implement, but hard for proxies and other SIP elements to recognize, as it would have to know about all number-context combinations in the world and track occasional changes. In addition, many of these numbers are being used for other services. For example, the emergency number in Paraguay (00) is also used to call the international operator in the United States. A number of countries, such as Italy, use 118 as an emergency number, but it also connects to directory assistance in Finland.

tel:sos This solution avoids name conflicts, but is not a valid "tel" [16] URI. It also only works if every outbound proxy knows how to route requests to a proxy that can reach emergency services since tel URIs. The SIP URI proposed here only requires a user's home domain to be appropriately configured.

sip:sos@domain Earlier work had defined a special user identifier, sos, within the caller's home domain in a SIP URI, for example, sip:sos@example.com. This approach had the advantage that dial plans in existing user agents could probably be converted to generate such a URI and that only the home proxy for the domain has to understand the user naming convention. However, it overloads the user part of the URI with specific semantics rather than being opaque, makes routing by the outbound proxy a special case that does not conform to normal SIP request-URI handling rules and is SIP-specific. The mechanism also does not extend readily to other services.

SIP URI user parameter: One could create a special URI, such as "aor-domain;user=sos". This avoids the name conflict problem, but requires mechanism-aware user agents that are capable of emitting this special URI. Also, the 'user' parameter is meant to describe the format of the user part of the SIP URI, which this usage does not do. Adding other parameters still leaves unclear what, if any, conventions should be used for the user and domain part of the URL. Neither solution is likely to be backward-compatible with existing clients.

Special domain: A special domain, such as "sip:fire@sos.int" could be used to identify emergency calls. This has similar properties as the "tel:sos" URI, except that it is indeed a valid URI. To make this usable, the special domain would have to be operational and point to an appropriate emergency services proxy. Having a single, if logical, emergency services proxy for the whole world seems to have undesirable scaling and administrative properties.

Appendix B. Acknowledgments

This document is based on discussions with Jonathan Rosenberg and benefitted from the comments of Leslie Daigle, Benja Fallenstein, Paul Kyzivat, Andrew Newton, Jonathan Rosenberg and Martin Thomas.

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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