Network Working Group Internet-Draft Expires: October 31, 2002 B. Campbell
dynamicsoft
May 2, 2002

Requirements for Binding Published Data to SIP Services draft-campbell-pub-bind-reqs-00

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of <u>Section 10 of RFC2026</u>.

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 October 31, 2002.

Copyright Notice

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

Abstract

A growing number of SIP based services depend on the idea that clients may publish service-related information to network elements. This information may then affect further processing of the service. Examples of such services include presence and CPL.

Multiple protocols may exists for the actual publication of such data. But regardless of the protocol, a client must know where to send it. This document describes requirements for a mechanism to inform clients where and how to publish service related information.

Expires October 31, 2002 [Page 1]

Table of Contents

<u>1</u> .	Introduction			<u>3</u>
<u>2</u> .	Scope of This Document			<u>3</u>
<u>3</u> .	Publication Mechanism Assumptions			<u>3</u>
<u>4</u> .	Requirements			<u>4</u>
<u>4.1</u>	1 Publication URI			<u>4</u>
<u>4.2</u>	<u>2</u> Publication Mechanism			<u>4</u>
<u>4.3</u>	<u>3</u> Address of Record			<u>4</u>
<u>4.4</u>	<u>4</u> Enumeration of Services			<u>4</u>
<u>4.5</u>	5 Lifetime of Publication URIs			<u>5</u>
<u>4.6</u>	<u>6</u> Communication of Publication URIs			<u>5</u>
<u>4.7</u>	7 Separate publication points			<u>5</u>
<u>4.8</u>	8 Publication URIs that are not easily guessable			<u>5</u>
<u>4.9</u>	9 Security	•	•	<u>5</u>
<u>5</u> .	Security Considerations			<u>5</u>
<u>6</u> .	Acknowledgements			<u>6</u>
	References			<u>6</u>
	Author's Address			<u>6</u>
	Full Copyright Statement			<u>7</u>

Expires October 31, 2002 [Page 2]

1. Introduction

A growing number of SIP [1] related services depend on client publication of some form of service data to network elements. This data is then either distributed as part of the service, or affect the processing of the service in some way. Examples include SIMPLE [3], where a client may publish a presence document to a remote Presence Agent which distributes the document to presence watchers. Another example is the Call Processing Language [2], where users may publish CPL scripts to SIP registrars.

The actual mechanisms for such publications are somewhat controversial. There is general agreement in the SIP and SIPPING working groups that the SIP REGISTER method is not adequate to the task. To a lesser degree, there has been agreement that the publication mechanism should not be SIP at all. There is, however, no consensus as to what the publication mechanism should be, or that a single mechanism makes sense in all situations.

Regardless of the publication mechanism or mechanisms chosen, clients will require a way to determine where to send such service data, and to bind that data to a particular service element or resource. This is essentially a matter of client configuration, and could conceivably handled through provisioning; however such a provisioning effort would likely become prohibitively complex for large installations.

2. Scope of This Document

This document discusses requirements for a service data binding mechanism. It does not propose a mechanism for realizing those requirements.

It does not attempt to address requirements or mechanism for the actual publication of service data. We make only a few very basic assumptions about such mechanisms.

3. Publication Mechanism Assumptions

We assume any service data publication mechanism or mechanisms will use a URI based addressing scheme. It is possible that the mechanism will not use URIs natively in its internal addressing scheme, but it will be possible to express such addresses in terms of URIs externally to the mechanism. For example, FTP does not natively use URIs, but there exists an FTP URI scheme that can be used to express the combination of an FTP server and a location in its file system.

We assume that more than one publication mechanism will exist, and

[Page 3]

that a single service may actually use multiple mechanisms.

We assume the existence of mechanisms which directly push service data to the publication points, as well as mechanisms where a client tells a service a URI where it can find the service data. (It is interesting to notice that indirect publication is also a type of binding operation and may have requirements similar to those in this document.)

<u>4</u>. Requirements

This section describes the requirements that are known at the time of this draft. Requirement discussion is still underway in the SIPPING working group. This document attempts to capture the requirements discussed so far.

4.1 Publication URI

The binding mechanism MUST allow a client to discover or infer a URI, or set of URIs, to which data may be published for a particular service. The mechanism MUST allow for any URI scheme, and MUST NOT make assumptions about how a specific publication mechanism interprets URIs.

4.2 Publication Mechanism

The binding mechanism MUST allow a client to discover or infer the mechanism, or set of mechanisms, to use to publish data to a particular publication URI.

4.3 Address of Record

The binding mechanism MUST allow clients to determine binding information based only on knowledge of an AoR. A client MAY allow provisioning of individual service publication bindings for an AoR. The binding mechanism MUST allow for multiple data-driven services for a single AoL, and MUST allow the client to distinguish one such service from another for publication purposes. (For example, the AoR "sip:alice@example.com" may have both presence and CPL services associated with it. A client must be able to avoid sending CPL to the presence service, or a presence document to the CPL services.)

4.4 Enumeration of Services

The binding mechanism SHOULD offer a way for a client to determine a list of all services for a given AoR for which it can publish data.

4.5 Lifetime of Publication URIs

The binding mechanism MUST allow a service element to manage the lifetime of a publication URI. It MUST allow long-lived publication URIs. It MAY also allow very short-lived publication URIs; for example, the URI may be single-use only.

4.6 Communication of Publication URIs

The binding mechanism MUST allow a service provider to communicate publication URIs to a client, directly or through a third party. For example the client might directly query a service element, or query a directory service. The mechanism MAY also provide a method for a client to infer publication URIs from an AoR without directly contacting the service elements, for example by using an algorithmic transformation, schema mapping, or the DNS.

<u>4.7</u> Separate publication points

The mechanism MUST NOT require all publication URIs for a given AoR to share the same host, address, or even domain. The mechanism MUST NOT require the publication point for a data-driven service to be colocated with the network element(s) that provide the service itself.

4.8 Publication URIs that are not easily guessable

The binding mechanism SHOULD allow the use of publication URIs that are not easily guessable.

4.9 Security

The binding mechanism MUST allow a client to authenticate the source of a publication URI. The mechanism MAY allow a publication URI provider to authenticate clients. The mechanism MUST allow a client to ensure that publication URIs have not been tampered with.

5. Security Considerations

We make the working assumption that service publication bindings are somewhat public knowledge. This document does not contain strong requirements for confidential transmission of bindings. On the contrary, algorithmic transformations, DNS approaches, etc. could cause such bindings to be completely public.

The service data to be published may, on the other hand, be very sensitive. Security of published data is in general an aspect of the publication method, and is out of scope for this document. But it is

May 2002

very important that a client can determine that a publication binding comes from a known source, and has not been tampered with. Otherwise, it would be possible for an attacker to provide a false binding, and trick a client into publishing potentially sensitive data to an unauthorized party.

<u>6</u>. Acknowledgements

Robert Sparks, Adam Roach, Sean Olson, Henning Schulzrinne, Jonathan Rosenburg, and James Undery all contributed substantially to the discussion about this subject.

References

- [1] Rosenberg, J. and H. Schulzrinne, "SIP: Session Initiation Protocol", <u>draft-ietf-sip-rfc2543bis-09</u> (work in progress), February 2002.
- [2] Lennox, J. and H. Schulzrinne, "Call Processing Language Framework and Requirements", <u>RFC 2824</u>, May 2000.
- [3] Rosenberg, J., "Session Initiation Protocol (SIP) Extensions for Presence", <u>draft-ietf-simple-presence-06</u> (work in progress), April 2002.

Author's Address

Ben Campbell dynamicsoft 5100 Tennyson Parkway Suite 1200 Plano, TX 75024

EMail: bcampbell@dynamicsoft.com

Expires October 31, 2002 [Page 6]

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.

Expires October 31, 2002 [Page 7]