

Address Resolution for Instant Messaging and Presence

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2004).

Abstract

Presence and instant messaging are defined in RFC 2778. The Common Profiles for Presence and Instant Messaging define two Universal Resource Identifier (URI) schemes: 'im' for INSTANT INBOXes and 'pres' for PRESENTITIES. This document provides guidance for locating the resources associated with URIs that employ these schemes.

Table of Contents

1. Introduction. . . . .	2
2. Terminology . . . . .	3
3. Address Resolution. . . . .	3
4. Domain Name Lookup. . . . .	4
5. Processing SRV RRs. . . . .	4
6. Processing Multiple Addresses . . . . .	5
7. Security Considerations . . . . .	5
8. IANA Considerations . . . . .	5
9. Contributors. . . . .	6
10. Normative References. . . . .	6
11. Author's Address. . . . .	7
12. Full Copyright Statement. . . . .	8

## 1. Introduction

Presence and instant messaging are defined in RFC 2778 [5]. The Common Profiles for Presence (CPP) [2] and Instant Messaging (CPIM) [1] define two Universal Resource Identifier (URI) schemes: 'im' for INSTANT INBOXes and 'pres' for PRESENTITIES. This document provides rules for locating the resources associated with URIs that employ these schemes via the Domain Name Service (DNS) [4]. These rules could no doubt be applied to the resolution of other URI schemes that are unrelated to instant messaging and presence.

CPIM and CPP both specify operations that have 'source' and 'destination' attributes. While only the semantics, not the syntax, of these attributes are defined by CPIM and CPP, many instant messaging and presence protocols today support the use of URIs to reflect the source and destination of their operations. The 'im' and 'pres' URI schemes allow such protocols to express the identities of the principals associated with a protocol exchange. When these operations pass through a CPIM or CPP gateway, these URIs could be relayed without modification, which has a number of desirable properties for the purposes of interoperability.

These URI schemes are also useful in cases where no CPIM/ CPP gatewaying will occur. If a particular principal's endpoint supports multiple instant messaging applications, for example, then a domain that identifies that host might use the sort of DNS records described in this document to provide greater compatibility with clients that support only one instant messaging protocol. A client would look up the corresponding record to the supported protocol, and learn how to contact the endpoint for that protocol. The principal in this instance would use an IM URI as their canonical address.

In some architectures, these URIs might also be used to locate a CPIM or CPP gateway that serves a particular domain. If a particular IM service provider wishes to operate CPIM/ CPP gateways in its own domain that map standard Internet protocols to an internal proprietary protocol, that gateway could be identified by an IM URI. In that case, the DNS records used to dereference the IM URI would serve a purpose similar to that of Mail Exchange (MX) records.

The system described in this document relies on the use of DNS service (SRV) [7] records and address (A) records.

## 2. Terminology

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in BCP 14, RFC 2119 [3] and indicate requirement levels for compliant implementations.

This memo makes use of the vocabulary defined in RFC 2778 [5]. Terms such as CLOSED, INSTANT INBOX, INSTANT MESSAGE, and OPEN are used in the same meaning as defined therein.

## 3. Address Resolution

A client determines the address of an appropriate system running a server, on behalf of the system referenced by the domain, by resolving the destination domain name that is part of the identifier to either an intermediate relay system or a final target system.

Only resolvable, fully-qualified, domain names (FQDNs) are permitted when domain names are used in an Instant Messaging (IM) URI (i.e., domain names that can be resolved to SRV [7] or A Resource Record (RR)).

The symbolic name used in the Service field of the SRV record is "\_im" for instant messaging and "\_pres" for presence (matching their respective URI schemes). However, the advertisement of these services in the DNS is incomplete if it does not include the protocol that will be used to instantiate the instant messaging or presence operations. Thus, the Protocol field of the SRV record contains an IANA-registered label corresponding to the underlying instant messaging or presence protocol being advertised (see Section 8 for more information on valid Protocol fields).

Taking the IM URI as a concrete example, a lookup is performed for SRVs for the target domain, a desired service (using the "\_im" Service label) and a desired IM transfer protocol. If the destination INSTANT INBOX is "im:fred@example.com", and the sender wishes to use an IM transfer protocol called "BIP" (and supposing "\_bip" were registered with IANA as a valid Protocol label for the IM Service), then a SRV lookup is performed for:

\_im.\_bip.example.com.

The same procedure is used for PRES URIs, with the "\_pres" Service label.

Some clients may support multiple instant messaging or presence protocols; in these cases they may make several such SRV queries, in an application-specific order, until they find one supported in common with the target domain.

#### 4. Domain Name Lookup

Once a client lexically identifies a domain to which instant messaging or presence operations will be delivered for processing, a DNS lookup **MUST** be performed to resolve the domain. The names **MUST** be fully-qualified domain names (FQDNs) -- mechanisms for inferring FQDNs from partial names or local aliases are a local matter.

The lookup first attempts to locate SRV RRs associated with the domain. If a canonical name (CNAME) RR is found instead, the resulting domain is processed as if it were the initial domain.

If one or more SRV RRs are found for a given domain, a sender **MUST NOT** utilize any A RRs associated with that domain unless they are located using the SRV RRs. If no SRV RRs are found, but an A RR is found, then the A RR is treated as if it was associated with an implicit SRV RR, with a preference of 0, pointing to that domain.

#### 5. Processing SRV RRs

The returned DNS RRs, if any, specify the next-hop server, which may be a protocol gateway or an endpoint.

Receiving systems that are registered for this DNS-based SRV resolution service list the transfer protocols by which they can be reached, either directly or through a translating gateway (using combinations of Service and Protocol labels as described above). The transfer-time choice of the IM transfer protocol to be used (and, therefore, to be resolved) is a local configuration option for each sending system.

Using this mechanism, seamless routing of IM traffic is possible, regardless of whether a gateway is necessary for interoperation. To achieve this transparency, a separate RR for a gateway must be present for each transfer protocol and domain pair that it serves.

## 6. Processing Multiple Addresses

When the lookup succeeds, the mapping can result in a list of alternative delivery addresses rather than a single address, because of multiple SRV records. For reliable operations, the client **MUST** be able to try each of the relevant addresses in this list in order, until a delivery attempt succeeds. However, there **MAY** also be a configurable limit on the number of alternate addresses that can be tried. In any case, the client **SHOULD** try at least two addresses.

Resolvers must follow the standard procedures in RFC 2782 [7] for handling the priority and weight fields of SRV records.

## 7. Security Considerations

The usage of IM and PRES URIs, and the DNS procedures in this document, introduce no security considerations beyond those described in the requirements for instant messaging and presence ([6]) and the SRV specification ([7]).

Subsequent registrations of Protocol labels for use with the "\_im" or "\_pres" Service labels **MUST**, however, explain any security considerations that arise from the use of the protocol in question with SRV.

## 8. IANA Considerations

This document reserves the use of "\_im" and "\_pres" Service labels. Since these relate to a service which may pass messages over a number of different message transports, they must be associated with a specific instant messaging or presence service.

In order to ensure that the association between "\_im" and "\_pres" and their respective underlying services are deterministic, the IANA has created two independent registries: the Instant Messaging SRV Protocol Label registry and the Presence SRV Protocol Label registry. For each registry, an entry shall consist of a label name and a pointer to a specification describing how the protocol named in the label uses SRV. Specifications should conform to the requirements listed in RFC 2434 [8] for "specification required".

Protocol labels compliant with this specification **MUST** begin with the underscore character "\_" and follow all other rules for SRV Protocol labels described in [7].

## 9. Contributors

Dave Crocker edited earlier versions of this document.

The following individuals made substantial textual contributions to this document:

Athanassios Diacakis (thanos.diacakis@openwave.com)

Florencio Mazzoldi (flo@networkprojects.com)

Christian Huitema (huitema@microsoft.com)

Graham Klyne (gk@ninebynine.org)

Jonathan Rosenberg (jdrosen@dynamicsoft.com)

Robert Sparks (rsparks@dynamicsoft.com)

Hiroyasu Sugano (suga@flab.fujitsu.co.jp)

## 10. Normative References

- [1] Peterson, J., "Common Profile for Instant Messaging (CPIM)", RFC 3860, August 2004.
- [2] Peterson, J., "Common Profile for Presence (CPP)", RFC 3859, August 2004.
- [3] Bradner, S., "Key words for use in RFCs to indicate requirement levels", BCP 14, RFC 2119, March 1997.
- [4] Mockapetris, P., "Domain Names - Concepts and Facilities", STD 13, RFC 1034, November 1987.
- [5] Day, M., Rosenberg, J., and H. Sugano, "A Model for Presence and Instant Messaging", RFC 2778, February 2000.
- [6] Day, M., Aggarwal, S., and J. Vincent, "Instant Messaging / Presence Protocol Requirements", RFC 2779, February 2000.
- [7] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for Specifying the Location of Services (SRV)", RFC 2782, February 2000.
- [8] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 2434, BCP 26, October 1998.

11. Author's Address

Jon Peterson  
NeuStar, Inc.  
1800 Sutter St  
Suite 570  
Concord, CA 94520  
US

Phone: +1 925/363-8720  
EMail: jon.peterson@neustar.biz

## 12. Full Copyright Statement

Copyright (C) The Internet Society (2004). 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.

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.

### Intellectual Property

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](mailto:ietf-ipr@ietf.org).

### Acknowledgement

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



