Internet Engineering Task Force (IETF)
Request for Comments: 7984
Updates: 3263
Category: Standards Track
ISSN: 2070-1721

O. Johansson
Edvina AB
G. Salgueiro
Cisco Systems
V. Gurbani
Bell Labs, Nokia Networks
D. Worley, Ed.
Ariadne

September 2016

Locating Session Initiation Protocol (SIP) Servers in a Dual-Stack IP Network

Abstract

RFC 3263 defines how a Session Initiation Protocol (SIP) implementation, given a SIP Uniform Resource Identifier (URI), should locate the next-hop SIP server using Domain Name System (DNS) procedures. As SIP networks increasingly transition from IPv4-only to dual-stack, a quality user experience must be ensured for dual-stack SIP implementations. This document updates the DNS procedures described in RFC 3263 for dual-stack SIP implementations in preparation for forthcoming specifications for applying "Happy Eyeballs" principles to SIP.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7984.

Table of Contents

1. Introduction

The Session Initiation Protocol (SIP) [RFC3261] and the additional documents that extended it provide support for both IPv4 and IPv6. However, this support does not fully extend to the highly hybridized environments that are characteristic of the transitional migratory phase from IPv4 to IPv6 networks. During this phase, many server and client implementations run on dual-stack hosts. In such environments, a dual-stack host will likely suffer greater connection delay, and by extension an inferior user experience, than an IPv4-only host. The need to remedy this diminished performance of dual-stack hosts led to the development of the "Happy Eyeballs" [RFC6555] algorithm, which has since been implemented in many protocols and applications.

This document updates the DNS lookup procedures of RFC 3263 [RFC3263] in preparation for the specification of the application of Happy Eyeballs to SIP. Happy Eyeballs will provide enhanced performance, and consequently enhanced user experience, in highly hybridized dual-stack SIP networks. The procedures described herein are such that a dual-stack client should look up both A and AAAA records in DNS and then select the best way to set up a network flow. The details of how the latter is done is considered out of scope for this document. See the Happy Eyeballs algorithm and implementation and design considerations in RFC 6555 [RFC6555] for more information about issues with setting up dual-stack network flows.

Section 4 of this document clarifies the interaction of [RFC3263] with [RFC6157] and [RFC6724].

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

RFC 3261 [RFC3261] defines additional terms used in this document that are specific to the SIP domain such as "proxy", "registrar", "redirect server", "user agent server" or "UAS", "user agent client" or "UAC", "back-to-back user agent" or "B2BUA", "dialog", "transaction", and "server transaction".

This document uses the term "SIP server" that is defined to include the following SIP entities: user agent server, registrar, redirect server, a SIP proxy in the role of user agent server, and a B2BUA in the role of a user agent server.

While this document focuses on the dual-stack situation described in RFC 6555 and other documents, concerning the migration from an IPv4-only network to a network supporting both IPv4 and IPv6, the techniques described can be used in other situations. Possible situations include when a device has multiple interfaces with distinct addressing characteristics and when additional IP address families are created in the future. This document uses the general term "dual-stack" to include all situations where the client has access to multiple communication methods that have distinct addressing characteristics.

The term "address records" means the DNS records that translate a domain name into addresses within the address family or families that the entity supports (as A records provide IPv4 addresses and AAAA records provide IPv6 addresses), regardless of whether the address family was defined before or after this document was approved.

3. DNS Procedures in a Dual-Stack Network

This specification introduces two normative DNS lookup procedures. These are designed to improve the performance of dual-stack clients in IPv4/IPv6 networks.

3.1. Dual-Stack SIP UA DNS Record Lookup Procedure

Once the transport protocol has been determined, the procedure for discovering an IP address if the TARGET is not a numeric IP address but the port is explicitly stated in the URI, is detailed in Section 4.2 of RFC 3263 [RFC3263]. The piece relevant to this discussion is:

> If the TARGET was not a numeric IP address, but a port is present in the URI, the client performs an A or AAAA record lookup of the domain name. The result will be a list of IP addresses, each of which can be contacted at the specific port from the URI and transport protocol determined previously.

Section 4.2 of RFC 3263 [RFC3263] also goes on to describe the procedure for discovering an IP address if the TARGET is not a numeric IP address, and no port is present in the URI. The piece relevant to this discussion is:

> If no SRV records were found, the client performs an A or AAAA record lookup of the domain name. The result will be a list of IP addresses, each of which can be contacted using the transport protocol determined previously, at the default port for that transport. Processing then proceeds as described above for an explicit port once the A or AAAA records have been looked up.

Happy Eyeballs [RFC6555] documents that looking up the "A or AAAA record" is not an effective practice for dual-stack clients and that it can add significant connection delay and greatly degrade user experience. Therefore, this document makes the following normative addendum to the DNS lookup procedures in Section 4.2 of RFC 3263 [RFC3263] for IPv4/IPv6 hybrid SIP networks and recommends it as a best practice for such dual-stack networks:

> The dual-stack client SHOULD look up address records for all address families that it supports for the domain name and add the resulting addresses to the list of IP addresses to be contacted. A client MUST be prepared for the existence of DNS resource records containing addresses in families that it does not support; if such records may be returned by the client’s DNS queries, such records MUST be ignored as unusable and the supported addresses used as specified herein.

3.2. Indicating Address Family Preference in DNS SRV Records

The Happy Eyeballs algorithm [RFC6555] is particularly effective for dual-stack HTTP client applications that have significant performance differences between their IPv4 and IPv6 network paths. This is because the client can initiate two TCP connections to the server, one using IPv4 and one using IPv6, and then use the connection that completes first. This works properly because the client can test each route by initiating a TCP connection, but simply opening a TCP connection to an HTTP server does not change the server's state; the client will send the HTTP request on only one connection.

Unfortunately, in common SIP situations, it is not possible to "race" simultaneous request attempts using two address families. If the SIP requests are transmitted as single UDP packets, sending two copies of the request to two different addresses risks having two copies of the request propagating through the SIP network at the same time. The difference between SIP and HTTP is that in SIP, the sender cannot test a route in a non-state-changing way.

(If two copies of the same request arrive at the destination client, the client SHOULD reject the second of them with a response code of 482 [RFC3261]. To convey information on why the request was rejected to the originator, the client can include a descriptive reason phrase, for example, "Merged Request". However, issuing the 482 response is not sufficient to prevent user-visible differences in behavior. A proxy that is upstream of the second request to arrive at the client may (almost immediately!) serially fork the second request to further destinations (e.g., the voicemail service for the destination client).)

In this common scenario, it is often necessary for a dual-stack client to indicate a preference for either IPv4 or IPv6. A service may use DNS SRV records to indicate such a preference for an address family. This way, a server with a high-latency and/or low-capacity IPv4 tunnel may indicate a preference for being contacted using IPv6. A server that wishes to do this can use the lowest SRV priority to publish host names that only resolve in IPv6 and the next priority with host names that resolve in both address families.

Note that host names that have addresses in only one address family are discouraged by [RFC6555]. Such special-purpose host names SHOULD be used only as described in this section, as targets of SRV records for an aggregate host name, where the aggregate host name ultimately resolves to addresses in all families supported by the client.

4. Clarification of Interaction with RFC 6724

Section 5 of [RFC6157] specifies that the addresses from the address records for a single target DNS name for a server's DNS name must be contacted in the order specified by the source and destination address selection algorithms defined in [RFC6724]. The set of addresses provided to a single invocation of the destination address selection algorithm MUST be the address records for the target DNS name in a single SRV record (or, if there are no SRV records, the DNS name in the URI or derived via NAPTR) -- the destination address selection algorithm MUST NOT reorder addresses derived from different SRV records. Typically, destination address selection is done by using the (relatively new) getaddrinfo() function to translate the target DNS name into a list of IPv4 and/or IPv6 addresses in the order in which they are to be contacted, as that function implements [RFC6724].

Thus, if SRV lookup on the server's DNS name is successful, the major ordering of the complete list of destination addresses is determined by the priority and weight fields of the SRV records (as specified in [RFC2782]), and the (minor) ordering among the destinations derived from the "target" field of a single SRV record is determined by [RFC6724].

For example, consider a server with DNS name example.com, with TCP transport specified. The relevant SRV records for example.com are:

```
_sip._tcp.example.com.  300 IN SRV 10 1 5060 sip-1.example.com.
_sip._tcp.example.com.  300 IN SRV 20 1 5060 sip-2.example.com.
```

The processing of [RFC2782] results in this ordered list of target domain names:

```
sip-1.example.com
sip-2.example.com
```

The address records for sip-1.example.com, as ordered by [RFC6724], are:

```
sip-1.example.com.  300 IN AAAA 2001:0db8:58:c02::face
sip-1.example.com.  300 IN AAAA 2001:0db8:c:a06::2:cafe
sip-1.example.com.  300 IN AAAA 2001:0db8:44:204::d1ce
sip-1.example.com.  300 IN A 192.0.2.45
sip-1.example.com.  300 IN A 203.0.113.109
sip-1.example.com.  300 IN A 198.51.100.24
```

And the address records for sip-2.example.com, as ordered by [RFC6724], are:

```
sip-2.example.com.  300 IN AAAA 2001:0db8:58:c02::dead
sip-2.example.com.  300 IN AAAA 2001:0db8:c:a06::2:beef
sip-2.example.com.  300 IN AAAA 2001:0db8:44:204::c0de
sip-2.example.com.  300 IN A 192.0.2.75
sip-2.example.com.  300 IN A 203.0.113.38
sip-2.example.com.  300 IN A 198.51.100.140
```

Thus, the complete list of destination addresses has this ordering:

```
2001:0db8:58:c02::face
2001:0db8:c:a06::2:cafe
2001:0db8:44:204::d1ce
192.0.2.45
203.0.113.109
198.51.100.24
2001:0db8:58:c02::dead
2001:0db8:c:a06::2:beef
2001:0db8:44:204::c0de
192.0.2.75
203.0.113.38
198.51.100.140
```

In particular, the destination addresses derived from sip-1.example.com and those derived from sip-2.example.com are not interleaved; [RFC6724] does not operate on the complete list. This would be true even if the two SRV records had the same priority and were (randomly) ordered based on their weights, as the address records of two target DNS names are never interleaved.

5. Security Considerations

This document introduces two new normative procedures to the existing DNS procedures used to locate SIP servers. A client may contact additional target addresses for a URI beyond those prescribed in [RFC3263], and/or it may contact target addresses in a different order than prescribed in [RFC3263]. Neither of these changes introduce any new security considerations because it has always been assumed that a client desiring to send to a URI may contact any of its targets that are listed in DNS.

The specific security vulnerabilities, attacks, and threat models of the various protocols discussed in this document (SIP, DNS, SRV records, Happy Eyeballs requirements and algorithm, etc.) are well documented in their respective specifications.

6. References

6.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, DOI 10.17487/RFC2782, February 2000, <http://www.rfc-editor.org/info/rfc2782>.

[RFC3263] Rosenberg, J. and H. Schulzrinne, "Session Initiation Protocol (SIP): Locating SIP Servers", RFC 3263, DOI 10.17487/RFC3263, June 2002, <http://www.rfc-editor.org/info/rfc3263>.

[RFC6157] Camarillo, G., El Malki, K., and V. Gurbani, "IPv6 Transition in the Session Initiation Protocol (SIP)", RFC 6157, DOI 10.17487/RFC6157, April 2011, <http://www.rfc-editor.org/info/rfc6157>.

[RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <http://www.rfc-editor.org/info/rfc6724>.

6.2. Informative References

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <http://www.rfc-editor.org/info/rfc3261>.

[RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", RFC 6555, DOI 10.17487/RFC6555, April 2012, <http://www.rfc-editor.org/info/rfc6555>.

Acknowledgments

The authors would like to acknowledge the support and contribution of the SIP Forum IPv6 Working Group. This document is based on a lot of tests and discussions at SIPit events, organized by the SIP Forum.

This document has benefited from the expertise and review feedback of many participants of the IETF DISPATCH and SIPCORE Working Group mailing lists as well as those on the SIP Forum IPv6 Task Group mailing list. The authors wish to specifically call out the efforts and express their gratitude for the detailed and thoughtful comments and corrections of Dan Wing, Brett Tate, Rifaat Shekh-Yusef, Carl Klatsky, Mary Barnes, Keith Drage, Cullen Jennings, Simon Perreault, Paul Kyzivat, Adam Roach, Richard Barnes, Ben Campbell, Stefan Winter, Spencer Dawkins, and Suresh Krishnan. Adam Roach devised the example in Section 4.

Authors' Addresses

Olle E. Johansson
Edvina AB
Runbovaegen 10
Sollentuna SE-192 48
Sweden

Email: oej@edvina.net

Gonzalo Salgueiro
Cisco Systems
7200-12 Kit Creek Road
Research Triangle Park, NC 27709
United States of America

Email: gsalguei@cisco.com

Vijay K. Gurbani
Bell Labs, Nokia Networks
1960 Lucent Lane
Rm 9C-533
Naperville, IL 60563
United States of America

Email: vkg@bell-labs.com

Dale R. Worley (editor)
Ariadne Internet Services
738 Main St.
Waltham, MA 02451
United States of America

Email: worley@ariadne.com