

MMUSIC
Internet-Draft
Intended status: Standards Track
Expires: January 12, 2014

T. Reddy
P. Patil
D. Wing
Cisco
July 11, 2013

Happy Eyeballs Extension for ICE
draft-reddy-mmusic-ice-happy-eyeballs-01

Abstract

This document specifies requirements for algorithms that make ICE connectivity checks more responsive by reducing delays in dual-stack host ICE connectivity checks when there is a path failure for the address family preferred by the application or by the operating system. As IPv6 is usually preferred, the procedures in this document helps avoid user-noticeable delays when the IPv6 path is broken or excessively slow.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on January 12, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Notational Conventions	2
3. Candidates Priority	2
4. Algorithm overview	3
4.1. Processing the Results	5
5. Indicating Happy-Eyeballs	6
6. IANA Considerations	6
7. Security Considerations	7
8. Acknowledgements	7
9. References	7
9.1. Normative References	7
9.2. Informative References	8
Authors' Addresses	8

1. Introduction

In situations where there are many IPv6 addresses, ICE [RFC5245] will prefer IPv6 candidates [RFC6724] and will attempt connectivity checks on all the IPv6 candidates before trying an IPv4 candidate. If the IPv6 path is broken, this fallback to IPv4 can consume a lot of time, harming user satisfaction of dual-stack devices.

This document describes an algorithm that makes ICE connectivity checks more responsive to failures of an address family by reordering the candidate pairs such that IPv6 and IPv4 candidates get a fair chance during connectivity checks. This document specifies requirements for any such algorithm, with the goals that the ICE agent need not be inordinately harmed with a simple reordering of the candidate pairs.

2. Notational Conventions

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 [RFC2119].

This note uses terminology defined in [RFC5245].

3. Candidates Priority

A prioritization formula is used by ICE [RFC5245] so that most preferred address pairs are tested first, and if a sufficiently good

pair is discovered, the tests can be stopped. With IPv6, addresses obtained from local network interfaces, called host candidates, are recommended as high-priority ones to be tested first since if they work, they provide usually the best path between the two hosts. The ICE specification recommends to use the rules defined in [RFC6724] as part of the prioritization formula for IPv6 host candidates and [I-D.keranen-mmusic-ice-address-selection] updates the ICE rules on how IPv6 host candidates are selected.

For dual-stack hosts the preference for IPv6 host candidates is higher than IPv4 host candidates based on precedence value of IP addresses described in [RFC6724]. IPv6 server reflexive candidates have higher precedence than IPv4 server reflexive candidate since NPTv6 is stateless and transport-agnostic.

```
(highest)  IPv6 Host Candidate
           IPv4 Host Candidate
           IPv6 Server Reflexive Candidate
           IPv4 Server Reflexive Candidate
           IPv6 Relayed Transport Candidate
(lowest)   IPv4 Relayed Transport Candidate
```

Figure 1: Candidate Preferences in decreasing order

By using the technique described in Section 4, if there are both IPv6 and IPv4 addresses in the check list, and the first 'N' candidates are of the same IP address family, then the highest-priority candidate pair of the other address family is promoted to position N in the check list thus making ICE connectivity checks more responsive to failures of an address family.

Note: The algorithm works even if the administrator changes the policy table to prefer IPv4 addresses over IPv6 addresses as defined in [RFC6724]

4. Algorithm overview

The Happy Eyeballs Extension for ICE algorithm proposes the following:

1. Indicate support for ICE Happy Eyeballs in the SDP offer if the end-point is dual-stack. The end-point will also include the position 'N' at which promotion is to occur.

2. After SDP offer/answer exchange if both end points support ICE Happy Extension for ICE algorithm the following steps are performed by the ICE agents after computing the candidate pair priority, ordering and pruning the pairs (section 5.7.2, 5.7.3 of [RFC5245])
 - a. If the first 'N' candidate pairs in the check list are of the same IP address family, then the highest-priority candidate pair of the other address family is promoted to position 'N' in the list.
 - b. Step b is repeated for candidate pairs that are next in the check list. This is continued until all candidate pairs of the preferred address family are exhausted.

The result of these steps is that after every consecutive 'N' candidate pairs of the preferred family, a candidate pair of the other family is inserted.

The following figure illustrates the result of the algorithm on candidate pairs:

Before Happy Eyeballs Extension for ICE algorithm :

```
-----
(highest)  IPv6 Host Candidate-1 pair
           IPv6 Host Candidate-2 pair
           IPv6 Host Candidate-3 pair
           IPv6 Host Candidate-4 pair
           IPv6 Host Candidate-5 pair
           IPv6 Host Candidate-6 pair
           IPv6 Host Candidate-7 pair
           IPv4 Host Candidate pair
           IPv6 Server Reflexive Candidate pair
           IPv4 Server Reflexive Candidate pair
           IPv6 Relayed Transport Candidate pair
(lowest)   IPv4 Relayed Transport Candidate pair
```

After Happy Eyeballs Extension for ICE algorithm :

```
-----
(highest)  IPv6 Host Candidate-1 pair
           IPv6 Host Candidate-2 pair
           IPv6 Host Candidate-3 pair
           IPv4 Host Candidate pair          ---> Promoted pair
           IPv6 Host Candidate-4 pair
           IPv6 Host Candidate-5 pair
           IPv6 Host Candidate-6 pair
           IPv4 Server Reflexive Candidate pair ---> Promoted pair
```

```

IPv6 Host Candidate-7 pair
IPv6 Server Reflexive Candidate pair
IPv6 Relayed Transport Candidate pair
(lowest) IPv4 Relayed Transport Candidate pair

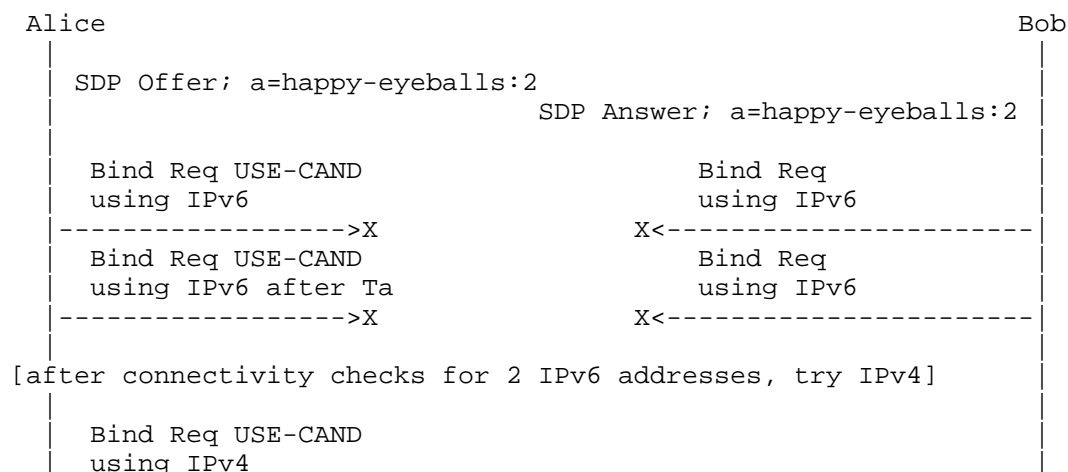
```

4.1. Processing the Results

If ICE connectivity checks using an IPv4 candidate is successful then ICE Agent will performs as usual "Discovering Peer Reflexive Candidates" (Section 7.1.3.2.1 of [RFC5245]), "Constructing a Valid Pair" (Section 7.1.3.2.2 of [RFC5245]), "Updating Pair States" (Section 7.1.3.2.3 of [RFC5245]), "Updating the Nominated Flag" (Section 7.1.3.2.4 of [RFC5245]).

If ICE connectivity checks using an IPv4 candidate is successful for each component of the media stream and connectivity checks using IPv6 candidates is not yet successful, the ICE endpoint will declare victory, conclude ICE for the media stream and start sending media using IPv4. However, it is also possible that ICE endpoint continues to perform ICE connectivity checks with IPv6 candidate pairs and if checks using higher-priority IPv6 candidate pair is successful then media stream can be moved to the IPv6 candidate pair. Continuing to perform connectivity checks can be useful for subsequent connections, to optimize which connectivity checks are tried first. Such optimization is out of scope of this document.

The following diagram shows the behaviour during the connectivity check when Alice calls Bob and Agent Alice is the controlling agent and uses the aggressive nomination algorithm. "USE-CAND" implies the presence of the USE-CANDIDATE attribute.






Figure 2: Happy Eyeballs Extension for ICE

5. Indicating Happy-Eyeballs

To indicate that Happy Eyeballs Extension for ICE algorithm defined in this document is used, the ICE offerer MUST include ice-options attribute with "happy-eyeballs:N" option identifier in the Session Description Protocol (SDP) [RFC4566] ICE offer, where N indicates the position at which promotion is to occur. If the ICE offer does not include this option tag, the answerer SHOULD NOT utilize the updated ICE algorithm defined in this document. If the offer included the option tag and the answerer supports this specification, the answerer SHOULD add the same option tag to the response and use the Happy Eyeballs Extension for ICE algorithm. If the ICE answer does not contain the option tag, the offerer SHOULD NOT use the updated ICE algorithm. Recommended value for 'N' is 3. As IPv6 becomes more prevalent, the value of 'N' can be increased as desired.

6. IANA Considerations

IANA is requested to register "happy-eyeballs" option identifier under the "ICE Options" [RFC6336] registry.

The required registration information is as follows:

Option identifier: happy-eyeballs

Contact: Tirumaleswar Reddy, tiredddy@cisco.com

Change control: IETF

Description: Existence of this option identifier indicates that Happy Eyeballs Extension for ICE algorithm is used.

Reference: RFCXXXX

[RFC editor: replace XXXX with the RFC number of this document]

7. Security Considerations

STUN connectivity check using MAC computed during key exchanged in the signaling channel provides message integrity and data origin authentication as described in section 2.5 of [RFC5245] apply to this use.

8. Acknowledgements

The authors would like to thank Bernard Aboba for his inputs.

9. References

9.1. Normative References

- [I-D.keranen-mmusic-ice-address-selection]
Keraenen, A. and J. Arkko, "Update on Candidate Address Selection for Interactive Connectivity Establishment (ICE)", draft-keranen-mmusic-ice-address-selection-01 (work in progress), July 2012.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.
- [RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, October 2008.

- [RFC5766] Mahy, R., Matthews, P., and J. Rosenberg, "Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN)", RFC 5766, April 2010.
- [RFC6336] Westerlund, M. and C. Perkins, "IANA Registry for Interactive Connectivity Establishment (ICE) Options", RFC 6336, July 2011.
- [RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012.

9.2. Informative References

- [RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, August 1999.

Authors' Addresses

Tirumaleswar Reddy
Cisco Systems, Inc.
Cessna Business Park, Varthur Hobli
Sarjapur Marathalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: tiredddy@cisco.com

Prashanth Patil
Cisco Systems, Inc.
Cessna Business Park, Varthur Hobli
Sarjapur Marthalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: praspati@cisco.com

Dan Wing
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, California 95134
USA

Email: dwing@cisco.com