

MMUSIC
Internet-Draft
Intended status: Standards Track
Expires: April 11, 2014

T. Reddy
P. Patil
D. Wing
Cisco
October 08, 2013

Happy Eyeballs Extension for ICE
draft-reddy-mmusic-ice-happy-eyeballs-03

Abstract

This document describes an algorithm that makes Interactive Connectivity Establishment (ICE) connectivity checks more responsive by reducing delays in dual-stack host ICE connectivity checks when there is a path failure for an address family preferred by the application or by the operating system. As IPv6 is usually preferred over IPv4, 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 April 11, 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. Candidate Priority	3
4. Algorithm overview	3
4.1. Processing the Results	4
5. IANA Considerations	5
6. Security Considerations	5
7. Acknowledgements	5
8. References	6
8.1. Normative References	6
8.2. Informative References	6
Authors' Addresses	7

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 causes ICE to perform terribly in cases where IPv6 doesn't work, which is still very commonplace. This document recommends an alternative prioritization for candidates that improves this situation with a goal that the ICE agent not be inordinately harmed by a simple reordering of the candidates.

This document describes an algorithm that makes ICE connectivity checks more responsive to failures of an address family by reordering the candidates such that IPv6 and IPv4 candidates get a fair chance during connectivity checks. This algorithm change is backward compatible with existing implementations, and does not require any changes other than to the selection of candidate priority.

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. Candidate Priority

By using the technique described in Section 4, if there are both IPv6 and IPv4 addresses candidates gathered, and the first 'N' candidates are of the same IP address family, then the highest-priority candidate of the other address family is promoted to position 'N+1' in the check list thus making ICE connectivity checks more responsive to failures of an address family. The algorithm ensures that there are no more than a fixed number of candidates of a given IP version in a single sequence.

Even if an administrator changes the policy table to prefer IPv4 addresses over IPv6 addresses as explained in [RFC6724], the IPv4 server-reflexive candidates will still have lower priority than IPv6 host candidates as per the "Recommended Formula" (section 4.1.2.1 of [RFC5245]) which is not desired. The Happy Eyeballs extension for ICE algorithm resolves the problem in this scenario as well by ensuring that IPv4 server-reflexive candidates are placed before IPv6 host candidates and thus ordering based on candidate types is no longer in effect.

4. Algorithm overview

The Happy Eyeballs Extension for ICE algorithm proposes the following steps after candidates are prioritized using the formula in section 4.1.2.1 of [RFC5245]:

- a. If the first 'N' candidates are of the same IP address family, then the highest-priority candidate of the other address family is promoted to position 'N+1' in the list.
- b. Step (a) is repeated for subsequent candidates in the list until all candidates of the preferred address family are exhausted.

The algorithm ensures that a long sequence of candidates belonging to the same address family is interleaved with candidates from an alternative IP version.

The following figure illustrates the result of the algorithm on candidates:

Before Happy Eyeballs Extension for ICE algorithm :

```
-----  
(highest)  IPv6 Host Candidate-1  
           IPv6 Host Candidate-2  
           IPv6 Host Candidate-3  
           IPv6 Host Candidate-4  
           IPv6 Host Candidate-5
```

```

IPv6 Host Candidate-6
IPv6 Host Candidate-7
IPv4 Host Candidate
IPv6 Server Reflexive Candidate
IPv4 Server Reflexive Candidate
IPv6 Relayed Transport Candidate
(lowest) IPv4 Relayed Transport Candidate

```

After Happy Eyeballs Extension for ICE algorithm :

```

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

```

4.1. Processing the Results

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 optimizations are 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 1: Happy Eyeballs Extension for ICE

5. IANA Considerations

None.

6. 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.

7. Acknowledgements

Authors would like to thank Bernard Aboba, Martin Thomson, Jonathan Lennox, Pal Martinson for their comments and review.

8. References

8.1. Normative References

- [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.

8.2. Informative 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.
- [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.
Bangalore
India

Email: praspati@cisco.com

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

Email: dwing@cisco.com