MMUSIC
Internet-Draft
Intended status: Standards Track
Expires: April 7, 2013

T. Reddy
P. Patil
D. Wing
Cisco

October 4, 2012

Happy Eyeballs Extension for ICE
draft-reddy-mmusic-ice-happy-eyeballs-00

Abstract

This document specifies requirements for algorithms that make ICE connectivity checks more aggressive to reduce delays in dual stack host 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-noticable delays wheen 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 7, 2013.

Copyright Notice

Table of Contents

1. Introduction 3
2. Notational Conventions 3
3. Candidates Priority 3
4. Algorithm overview 4
4.1. Processing the Results 5
5. Relayed Candidates 7
6. Setting Te, Tr and MAX_PAIRS_HAPPYEYE_STAGE 8
7. IANA Considerations 8
8. Security Considerations 8
9. References 8
9.1. Normative References 8
9.2. Informative References 9
Authors' Addresses 9

1. Introduction

In situations where there are many IPv6 addresses, ICE [RFC5245] will prefer IPv6 [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 performing connectivity checks with both IPv6 and IPv4 candidates in parallel if IPv6 connectivity checks have not yet succeeded. This document specifies requirements for any such algorithm, with the goals that the ICE agent need not be inordinately harmed with a simple parallelisation of IPv6 and IPv4 connectivity checks and ensuring that the priority of precedence defined in [RFC6724] be honored.

For either of the address families, there is also a very realistic chance that connectivity checks for relayed candidates will always work. There are scenarios where firewalls block connectivity checks for Host/Server Reflexive candidates or for IPv4 or for IPv6. This document also proposes an optimization where connectivity checks with relayed checks are performed earlier than usual if connectivity checks using other candidates do not succeed.

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 in Section 4 IPv6 candidate pairs will be tested first as usual, but if connectivity checks are not successful after a certain period of time, the algorithm will become more aggressive and connectivity checks using IPv6/IPv4 host/ server-reflexive candidates will be performed simultaneously. If connectivity checks with IPv6 candidate pairs do not yield any successful result then ICE endpoints can immediately start sending media using IPv4 host/server-reflexive candidates.

Note: [RFC6724] permits administrator to change the policy table to prefer IPv4 addresses over IPv6 addresses in which case the algorithm described in the next section is reversed.

4. Algorithm overview

The Happy Eyeballs Extension for ICE is governed by a timer (Te) that is started just before carrying out the ICE connectivity checks for each check list under the following conditions:

1. when the candidates pairs include IPv6 and IPv4 addresses

2. list of IPv6 candidate pairs is higher than a configured threshold (MAX_PAIRS_HAPPYEYE_STAGE_I). [RFC5245] recommends a limit of 100 for the candidate pairs.

When the timer (Te) fires, if the connectivity check using IPv6 candidate pairs are not yet successful and if the number of IPv6 candidate pairs with remote candidates of type host in the check list that are in Waiting and Frozen state are non-zero, the ICE agent performs the following Happy Eyeball steps in parallel with the regular ICE Ordinary checks:

- Find the highest priority pair in the checklist that is in the Waiting state with candidate address family being IPv4 and remote candidate of type host. If there are no remote IPv6 candidates of type server-reflexive then IPv4 remote candidates of type server-reflexive will be added to the search.

    1. If there is such a pair then perform ICE connectivity check on this pair and set the state of the candidate pair to In-Progress.

    2. If there is no such pair find the highest priority pair in the checklist that is in the Frozen state with candidate address family being IPv4 and remote candidate of type host candidate. If there are no remote IPv6 candidates of type server-reflexive then IPv4 remote candidates of type server-reflexive will be added to the search. If there is such pair in Frozen state then unfreeze the pair, perform connectivity check on this pair and set the state of the candidate pair to In-Progress.

- The above mentioned steps will be followed every Ta milliseconds and stopped when any of the below conditions are met:

1. All IPv6 candidate pairs with remote candidates of type host in the check list are in any of the following states Succeeded, In-Progress or Failed states. The parallel activity is not required beyond this point because the regular ICE algorithm will itself pick up IPv4 candidate pairs not yet tested.

2. All IPv4 candidate pairs with remote candidates of type host/ server reflexive are in any of the following states Succeeded, In-Progress or Failed states.

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.

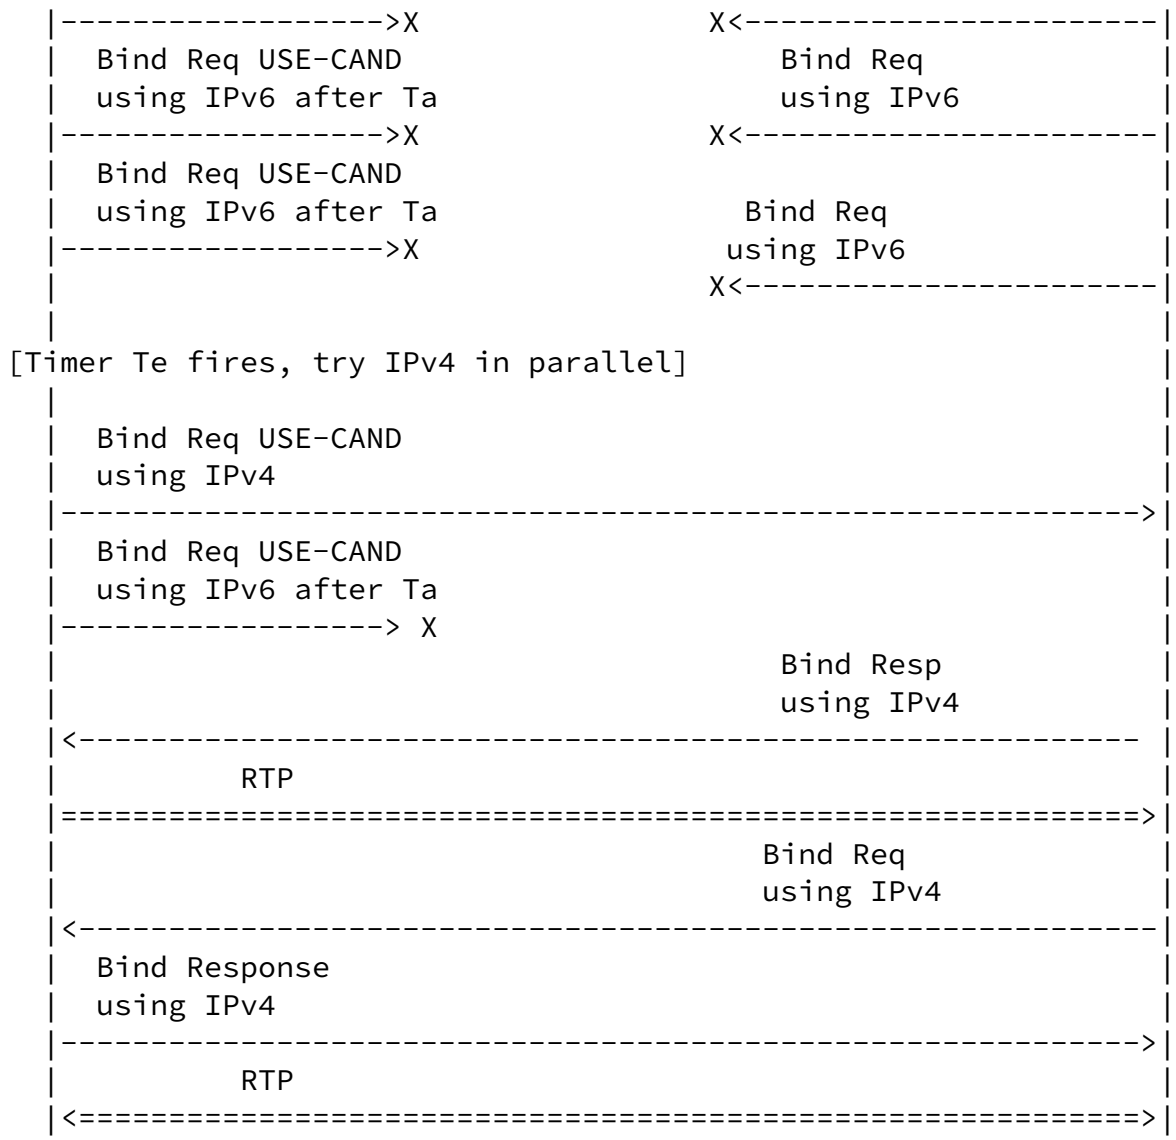

Figure 2: Happy Eyeballs Extension for ICE

5. Relayed Candidates

The optimization proposes doing connectivity checks with relayed candidates in parallel with other candidates. The algorithm does not make a distinction between IPv6/IPv4 relayed candidates and will choose the existing order among relayed candidate pair defined by ICE. If ICE connectivity check is successful using a relayed candidate from either of the IP address families, the ICE agent can

stop connectivity checks for other relayed candidates.

This part of the Happy Eyeballs Extension for ICE is governed by a timer (Tr) that is started just before carrying out the ICE connectivity checks for each check list under the following conditions:

1. when the candidates pairs include IPv6 and IPv4 relayed addresses

2. list of candidate pairs is higher than a configured threshold (MAX_PAIRS_HAPPYEYE_STAGE_I).

When the timer (Tr) fires, If no ICE connectivity checks are successful as yet and if ICE Connectivity checks using IPv6 and IPv4 local relayed candidates have not yet been attempted then the following steps will be started by the ICE agent in parallel with other connectivity checks:

o Find the highest priority pair in the checklist that is in the Waiting state with local candidate of type relayed.

1. If there is such a pair then perform ICE connectivity check on this pair and set the state of the candidate pair to In-Progress.

2. If there is no such pair find the highest priority pair in the checklist that is in the Frozen state with local candidate of type relayed. If there is such pair in Frozen state then unfreeze the pair, perform connectivity check on this pair and set the state of the candidate pair to In-Progress.

If ICE connectivity checks using relayed candidate is successful then ICE Agent will performs as usual "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 local relayed candidates is successful for each component of the media stream and connectivity checks using higher priority candidate pairs has not yet succeeded then conclude ICE for the media stream and proceed to send media using local relayed candidate.

However ICE connectivity checks MUST be continued and if the check

succeeds for a pair whose priority is higher than the previously selected candidate pair then media session will be moved to this pair. Hence media will only be sent briefly on TURN relays. Additional TURN server load is created due to this recommendations, especially when connectivity check using IPv6/IPv4 host/ server-reflexive candidates are not completing quickly and the side affect could be that RTP receivers will receive packets out of order during switchover.

6. Setting Te, Tr and MAX_PAIRS_HAPPYEYE_STAGE

The value of Ta, Tr, MAX_PAIRS_HAPPYEYE_STAGE_I, MAX_PAIRS_HAPPYEYE_STAGE_II and SHOULD be configurable, and SHOULD have a default of:

```
Te : 150ms
Tr : 500ms
MAX_PAIRS_HAPPYEYE_STAGE_I  : 12
MAX_PAIRS_HAPPYEYE_STAGE_II :  6
```

Figure 3: Default Values

7. IANA Considerations

None.

8. 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.

9. References

9.1. Normative References

[I-D.keranen-mmusic-ice-address-selection]
Keranen, 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.

[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.

[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: tireddy@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