

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2015

J. Uberti
Google
J. Lennox
Vidyo
March 09, 2015

Improvements to ICE Candidate Nomination
draft-uberti-mmusic-nombis-00

Abstract

This document makes recommendations for simplifying and improving the procedures for candidate nomination in Interactive Connectivity Establishment (ICE).

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 September 10, 2015.

Copyright Notice

Copyright (c) 2015 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.	Terminology	3
3.	Goals and Requirements	3
3.1.	Minimize Call Setup Latency	3
3.2.	Allow Controlling Endpoint to Make Dynamic Decisions	3
3.3.	Allow Selected Pair Change At Any Time Without Signaling	4
3.4.	Allow Continuous Addition of Candidates	4
3.5.	Maintain Backwards Compatibility	4
3.6.	Minimize Complexity Increase	5
4.	Deprecating Aggressive Nomination	5
4.1.	Overview	5
4.2.	Operation	5
4.3.	Backwards Compatibility	6
4.4.	Examples	6
5.	Introducing Continuous Nomination	7
5.1.	Overview	7
5.2.	Operation	8
5.3.	Backwards Compatibility	9
5.4.	Examples	9
5.4.1.	Switching Between Pairs Based on RTT	9
5.4.2.	Switching To A New TURN Server	9
5.4.3.	Switching From WLAN to WWAN	10
6.	Security Considerations	10
7.	IANA Considerations	10
8.	Acknowledgements	10
9.	References	10
9.1.	Normative References	10
9.2.	Informative References	11
Appendix A.	Change log	11
Authors' Addresses	12

1. Introduction

Interactive Connectivity Establishment (ICE) attempts to find the 'best' path for connectivity between two peers; in ICE parlance, these paths are known as 'candidate pairs'. During the ICE process, one endpoint, known as the 'controlling' endpoint, selects a candidate pair as the best pair; this action is known as nomination. ICE supports two different mechanisms for performing nomination, known as Regular Nomination, and Aggressive Nomination.

However, each of these modes have flaws that restrict their usefulness. Regular Nomination, as currently speced, requires a best pair to be chosen before media transmission can start, causing unnecessary call setup delay. Aggressive Nomination, while avoiding this delay, gives the controlling endpoint much less discretion into

which candidate pair is chosen, preventing it from making decisions based on dynamic factors such as RTT or loss rate. Needless to say, the presence of both modes also adds nontrivial complexity.

Lastly, ICE is currently defined as a finite process, where the decision on the optimal candidate pair is made during call setup and infrequently (if ever) changed. While this may be acceptable for endpoints with static network configurations, it fails to meet the needs of mobile endpoints, who may need to seamlessly move between networks, or be connected to multiple networks simultaneously. In these cases, the controlling endpoint may want to maintain multiple potential candidate pairs, and make dynamic decisions to switch between them as conditions change.

To address these challenges, this document makes two proposals for refactoring ICE nomination - merging Regular and Aggressive Nomination, and introducing a new mode, known as Continuous Nomination. This makes ICE substantially more flexible without increasing complexity.

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 [RFC2119].

3. Goals and Requirements

The goals for improved ICE nomination are enumerated below.

3.1. Minimize Call Setup Latency

Modern ICE agents will often have multiple network interfaces and multiple servers from which to obtain ICE candidates. While some ICE checks may succeed quickly, finishing the entire set of checks can easily take multiple seconds; this concern is discussed in [RFC5245], Section 8.1.1.1. As a result, ICE endpoints MUST be able to start transmitting media immediately upon a successful ICE check, and MUST retain the ability to switch if a better candidate pair becomes available later.

3.2. Allow Controlling Endpoint to Make Dynamic Decisions

While an ICE endpoint will assign various priority values to its ICE candidates, these priorities are static and can only be based on a priori knowledge; the shortcomings of this approach are discussed in the first paragraph of Section 2.6 in [RFC5245]. To properly make choices in multi-network and multi-server scenarios, the controlling

endpoint MUST be able to make dynamic decisions about the selected candidate pair based on observed network performance. For example, RTT could be used to evaluate which TURN servers to use, as described in [I-D.williams-peer-redirect]. To ensure symmetric flows, this implies that the controlling endpoint MUST be able to communicate its choice to the controlled side.

3.3. Allow Selected Pair Change At Any Time Without Signaling

Expanding on the requirement above, the need to make dynamic decisions is not limited to call setup. A multihomed endpoint may need to switch interfaces based on mobility considerations, or a robust endpoint may want to keep multiple network paths warm and switch immediately if connectivity is interrupted on one of them. As the signaling channel may be affected by the event necessitating the switch, this implies that the controlling endpoint MUST be able to change the selected pair and indicate this to the remote side without signaling. The need for this functionality has been stated in [I-D.wing-mmusic-ice-mobility] and [I-D.singh-avtcore-mprtp].

The rules in [RFC5245] ensure that the controlled endpoint keeps its candidate needed for the selected pair alive. However, in order for alternate pairs to remain available, the controlled endpoint must keep the associated candidates alive as well, following the procedures outlined in [RFC5245], Section 4.1.1.4. This implies that the controlling endpoint MUST have some way to indicate to the controlled side that specific candidates are to be kept alive.

3.4. Allow Continuous Addition of Candidates

In certain network mobility scenarios, networks may come up and down while the call is active. In order to allow candidates gathered on newly available networks to be used for the selected pair or backup pairs, the endpoint MUST be able to gather candidates on these networks and communicate them to the remote side. While this could be done using an ICE restart, as described in [RFC5245], Section 9.1, the ICE restart may have unintended consequences, such as causing the remote side to regather all candidates. Instead, it would be best if the new candidates could be trickled, as discussed in [I-D.ietf-mmusic-trickle-ice], but even after ICE processing has completed.

3.5. Maintain Backwards Compatibility

To prevent interoperability problems, ICE endpoints that support the functionality listed above MUST still maintain [RFC5245] compliance when interacting with existing endpoints. However, the ideal

solution SHOULD allow some improvements to occur when only the controlling side supports the new functionality.

3.6. Minimize Complexity Increase

Increased functionality typically leads to increased complexity, which leads to more edge cases, and more implementation bugs. This suggests that in addition to proposing new ICE functionality, the ideal solution SHOULD deprecate superfluous functionality.

4. Deprecating Aggressive Nomination

4.1. Overview

The main benefits of Regular Nomination are that the controlling side can dynamically choose which candidate pair to use, and a clear signal when the nomination process has completed, via the presence of the USE-CANDIDATE flag in a Binding Request. The main benefit of Aggressive Nomination is that it is only necessary to send a single Binding Request before starting the transmission of media, reducing setup latency. Why don't we have both?

By preserving the dynamic behavior of Regular Nomination, but allowing media transmission to start upon a single successful connectivity check, as in Aggressive Nomination, the requirements of Section 3.1 and Section 3.2 can be met, while meeting the compatibility requirement from Section 3.5 and, since Aggressive Nomination is no longer needed, the complexity requirement from Section 3.6.

4.2. Operation

Since media may be transmitted as soon as all components have a valid pair, as indicated in [RFC5245], Page 69, an ICE Agent can begin transmitting media as soon as this occurs, even if it has not sent a Binding Request with USE-CANDIDATE.

This pair can change as more pairs are added to the Valid list on the controlling side. When nomination completes, and a final pair is selected, this is communicated to the controlled side via the typical Binding Request with USE-CANDIDATE.

On the controlled side, the same process can occur, with the ICE Agent transmitting media as soon as a valid pair exists. To encourage use of symmetric RTP, the controlled ICE Agent SHOULD use the same candidate pair on which it received media from the controlling side. [Doesn't need to be secure media, since the

controlling side will finalize this preference through USE-CANDIDATE shortly.]

As this is legal ICE behavior, no negotiation of this mechanism should be needed. In the event the receiver drops any packets that arrive before a Binding Request with USE-CANDIDATE set, this will simply lead to brief media clipping and will resolve itself once nomination completes.

4.3. Backwards Compatibility

When acting in the controlled role, new implementations MUST NOT use Aggressive Nomination.

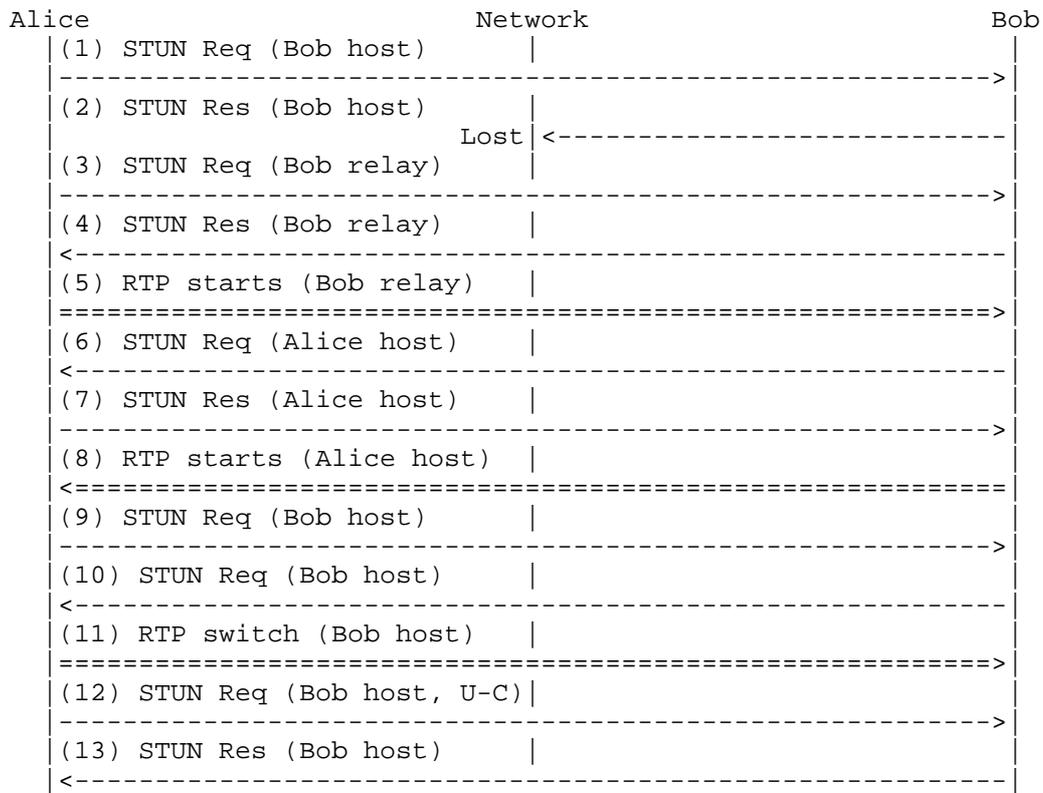
When acting in the controlled role, and the controlling side is using Aggressive Nomination (e.g. sending USE-CANDIDATE in its initial Binding Requests), the standard PRIORITY-based mechanism outlined in [RFC5245], Section 8.1.1.2 should be used to determine the reverse media path.

Note that if implementations would prefer to just avoid Aggressive Nomination altogether, they MAY indicate some TBD pseudo-option in the ice-options attribute. Because compliant implementations MUST NOT use Aggressive Nomination if an unknown ICE option is encountered, this effectively prohibits the use of Aggressive Nomination. [N.B. this could be the ice-options:continuous option described below]

4.4. Examples

An example call setup using Regular Nomination as described above is shown here. Alice is in the controlling role, and Bob is in the controlled role; Alice has a single host candidate and Bob has both host and relay candidates.

Alice's initial check to Bob's host candidate fails, but the check to his relay candidate succeeds, so Alice starts transmitting media on her host-relay pair. Bob's initial check from his host candidate to Alice's host candidate succeeds, so he starts transmitting media over this host-host pair to Alice. However, when Alice's host check is later retransmitted, it succeeds, and Alice determines that the host-host pair has a better RTT than the host-relay pair, so she cuts media over to use the host-host pair. Eventually, Alice concludes Regular Nomination by sending a final check to Bob with the USE-CANDIDATE flag set. If Bob had selected a different pair to use than Alice, this action would have forced Bob to use the same pair.



5. Introducing Continuous Nomination

5.1. Overview

As discussed above, in mobile environments there can be multiple possible valid candidate pairs, and these can change at various points in the call, as new interfaces go up and down, signal strength for wireless interfaces changes, and new relay servers are discovered.

However, under 5245 rules, once a candidate pair is selected and confirmed, via USE-CANDIDATE, nomination has completed and cannot be restarted without performing an ICE restart. This is overly complex in many cases, and especially problematic in some specific ones, namely a wifi-cellular handover, where the signaling path for communicating an ICE restart may be impacted by the handover.

To address this situation, this section introduces the concept of "continuous nomination", where the controlling ICE endpoint can adjust the selected candidate pair at any time. By allowing ICE

processing to occur continuously during a call, rather than just at call setup, the requirements expressed in Section 3.3 and Section 3.4 can be met.

5.2. Operation

Under continuous nomination, ICE never concludes; new candidates can always be trickled, and a new candidate pair can be selected by the controlling side at any time.

When selecting a new candidate pair, the controlling side informs the controlled side of the chosen path by sending a new Binding Request with a USE-CANDIDATE attribute. The decision about which candidate pair to use is fully dynamic; the controlling side can use metrics such as RTT or loss rate to change the selected pair at any time. If Binding Requests need to be sent for any other reason, such as consent checks [I-D.ietf-rtcweb-stun-consent-freshness], any checks sent on the selected pair MUST include a USE-CANDIDATE attribute.

Upon receipt of a Binding Request with USE-CANDIDATE, the controlled side MUST switch its media path to the candidate pair on which the Binding Request was received.

During continuous nomination, the controlling side may still elect to prune certain candidate pairs; for example, an implementation may choose to drop relay candidates once a successful connection has been established. The controlled side, however, should follow the controlling side's lead in terms of deciding whether any pairs should be pruned. [TODO: should the controlled side have any say in the matter, e.g. to eliminate certain candidates?] The controlling ICE Agent informs the remote side of its preferences by continuing to send Binding Requests to the remote side on each candidate pair that it wants to retain. The controlled ICE Agent SHOULD prune any candidate pairs that have not received a Binding Request in N seconds (30?), and SHOULD NOT keep alive any candidates that are not associated with a live candidate pair. [TODO: decide if this implicit timeout approach is correct, or if we should have some sort of approach similar to TURN LIFETIME indicating when a pair should be GCed, with LIFETIME==0 indicating immediate GC.] One side benefit of doing this is that the continuous exchange of Binding Requests across all candidate pairs allows the RTT and loss rate for each to be reliably determined and kept up to date.

If the endpoints have negotiated Trickle ICE support [I-D.ietf-mmusic-trickle-ice], and new candidates become available on either side, the endpoint may send these candidates to the remote side using the existing Trickle ICE mechanisms. Once all of the new candidates have been transmitted, the endpoint MUST send an end-of-

candidates messages, which indicates that no more candidates will be sent in the near future.

At any point, either side may perform an ICE restart, which will result in both sides gathering new ICE candidates, starting a new continuous nomination sequence, and upon successful completion, discarding all candidates from the previous nomination sequence.

5.3. Backwards Compatibility

Since standard ICE implementations may not expect the selected pair to change after a USE-CANDIDATE attribute is received, support for continuous nomination is explicitly indicated via a new "continuous" value for ice-options. If the remote side does not support the "continuous" option, the controlling side MUST fall back to Regular Nomination, as specified in [RFC5245], Section 8.1.1.

5.4. Examples

5.4.1. Switching Between Pairs Based on RTT

Alice and Bob have set up a call using ICE and have established multiple valid pairs. The currently selected pair is for a peer-to-peer route, as it had the highest initial priority value. However, they have also kept alive a selected pair that goes through their TURN servers. At a certain point, Alice detects, via the connectivity checks that she continues to do on the relayed pair, that it actually has a better RTT than the peer-to-peer path. She then decides to switch media over to this path.

As mentioned above, this is easily handled by Alice immediately switching her media to the relayed path; future STUN checks on this path also include the USE-CANDIDATE attribute.

5.4.2. Switching To A New TURN Server

Alice and Bob have set up a call using ICE, and are currently sending their media through Alice's TURN server. At a certain point, Alice's application discovers a new TURN server that it thinks might provide a better path for this call.

Alice gathers new candidates from this TURN server, and trickles them to Bob. They perform connectivity checks using these candidates, and Alice determines that the RTT when going through this TURN server is better than the RTT of the current relayed path.

As in the previous example, this is easily handled by Alice switching media to the new path, along with sending USE-CANDIDATE. If the old

path is no longer needed, Alice can destroy the allocation on the old TURN server, and Bob will stop checking it when it stops working.

5.4.3. Switching From WLAN to WWAN

Alice and Bob have set up a call using ICE, and are currently exchanging their media directly via a peer-to-peer path. Alice is on a mobile device, with both wifi and cellular interfaces, but for power reasons, candidates have only been gathered on the wifi interface. At a certain point, Alice leaves her home while the call is active.

In response to the decreasing wifi signal strength, Alice starts to collect candidates on the cellular interface, and trickles them to Bob. They perform connectivity checks using these candidates, and, because of the low wifi signal strength, these candidates are preferred over the existing selected pair.

As in the previous examples, Alice can easily switch media to the new selected pair. When Alice walks completely out of wifi range, and the wifi interface goes down, the wifi candidates are pruned, and any valid pairs on Bob's side that use those candidates will time out and be pruned as well.

6. Security Considerations

TODO

7. IANA Considerations

A new ICE option "continuous" has been [will be] registered in the "ICE Options" registry created by [RFC6336].

8. Acknowledgements

Several people provided significant input into this document, including Martin Thomson, Brandon Williams, and Dan Wing. Emil Ivov also provided several of the examples for continuous nomination.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.
- [RFC6336] Westerlund, M. and C. Perkins, "IANA Registry for Interactive Connectivity Establishment (ICE) Options", RFC 6336, July 2011.

9.2. Informative References

- [I-D.ietf-mmusic-trickle-ice]
Ivov, E., Rescorla, E., and J. Uberti, "Trickle ICE: Incremental Provisioning of Candidates for the Interactive Connectivity Establishment (ICE) Protocol", draft-ietf-mmusic-trickle-ice-02 (work in progress), January 2015.
- [I-D.ietf-rtcweb-stun-consent-freshness]
Perumal, M., Wing, D., R, R., Reddy, T., and M. Thomson, "STUN Usage for Consent Freshness", draft-ietf-rtcweb-stun-consent-freshness-11 (work in progress), December 2014.
- [I-D.singh-avtcore-mprtp]
Singh, V., Karkkainen, T., Ott, J., Ahsan, S., and L. Eggert, "Multipath RTP (MPRTP)", draft-singh-avtcore-mprtp-10 (work in progress), November 2014.
- [I-D.williams-peer-redirect]
Williams, B. and T. Reddy, "Peer-specific Redirection for Traversal Using Relays around NAT (TURN)", draft-williams-peer-redirect-03 (work in progress), December 2014.
- [I-D.wing-mmusic-ice-mobility]
Wing, D., Reddy, T., Patil, P., and P. Martinsen, "Mobility with ICE (MICE)", draft-wing-mmusic-ice-mobility-07 (work in progress), June 2014.

Appendix A. Change log

Changes in draft -00:

- o Initial version, from mailing list discussion post-IETF 90.

Authors' Addresses

Justin Uberti
Google
747 6th Ave S
Kirkland, WA 98033
USA

Email: justin@uberti.name

Jonathan Lennox
Vidyo
433 Hackensack Avenue
Hackensack, NJ 07601
USA

Email: jonathan@vidyo.com