

MMUSIC
Internet-Draft
Intended status: Informational
Expires: April 13, 2014

T. Reddy
D. Wing
B. VerSteeg
R. Penno
Cisco
V. Singh
Aalto University
October 10, 2013

Improving ICE Interface Selection Using Port Control Protocol (PCP) Flow
Extension
draft-reddy-mmusic-ice-best-interface-pcp-00

Abstract

A host with multiple interfaces needs to choose the best interface for communication. Oftentimes, this decision is based on a static configuration and does not consider the link characteristics of that interface, which may affect the user experience.

This document describes a mechanism for an endpoint to query the link characteristics from the access router (the router at the other end of the endpoint's access link) using a Port Control Protocol (PCP) Flow Extension. This information influences endpoint's Interactive Connectivity Establishment (ICE) candidate selection algorithm.

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 13, 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	3
2. Notational Conventions	3
3. Algorithm overview	3
3.1. Changed Link Quality	4
4. Multiple Interfaces	5
4.1. Multiple Interfaces for media streams	5
4.2. Availability of New Interfaces	5
5. IANA Considerations	6
6. Security Considerations	6
7. Acknowledgements	6
8. References	6
8.1. Normative References	6
8.2. Informative References	6
Appendix A.	7
A.1. Delay Factor in Discovering the Link Characteristics	7
Authors' Addresses	7

1. Introduction

ICE [RFC5245] uses a prioritization formula to perform connectivity checks, in which the most preferred address pairs are tested first and when a sufficiently good pair is discovered, the ICE connectivity tests are stopped. ICE prefers address pairs in the following order: transport address directly attached to the endpoint's network interface (host candidate), transport address on the public side of a NAT (server reflexive candidate), and finally, transport address that are allocated on a media relay (relayed candidate). This approach works well for an endpoint with a single interface, but is too simplistic for endpoints with multiple interfaces. The network interfaces may have different link characteristics, but that will not be known without the awareness of the upstream and downstream characteristics of the access link.

In this document, an ICE agent [RFC5245] uses PCP Flow Extension [I-D.wing-pcp-flowdata] to determine the link characteristics of the host's interfaces, which influence the ICE candidate priority.

As this document explains the interworking of ICE and Port Control Protocol (PCP) Flow Extensions, it is beneficial to first read the Overview of ICE (Section 2 of ICE [RFC5245]) and PCP Flowdata Option [I-D.wing-pcp-flowdata]. Additionally, PCP for WebRTC [I-D.penno-rtcweb-pcp] describes the problems with traversing NATs and firewalls, current techniques used to solve them and the PCP solution in these scenarios.

2. Notational Conventions

This note uses terminology defined in ICE [RFC5245] and PCP [RFC6887].

3. Algorithm overview

The proposed algorithm is backward compatible with existing implementations, and does not require any changes other than to the selection of candidate priority.

When an endpoint first joins a network, it determines if the network supports PCP Flow Extensions by following the procedures described in [I-D.wing-pcp-flowdata]. Basically, the endpoint sends a PCP Flow Extension probe packet, the response to which provides coarse information on the link capabilities. After confirming that PCP Flow Extensions are supported on that network interface, the ICE agent can use PCP Flow Extensions on that interface (rather than STUN).

When a media session needs to be established, and the user and operator controlled policies on an endpoint permit more than one interface for a media session, the ICE agent uses PCP Flow Extensions to (a) obtain a mapping from its NAT or firewall and (b) determine the characteristics of the link. After receiving the PCP Flow Extension responses from its various interfaces, the ICE agent sorts the ICE candidates according to the link capacity characteristics. ICE candidates from the interface which best fulfills the desired flow characteristics is assigned the highest priority and the best suited interface should be used to communicate with the TURN server to learn the relayed candidate address.

The ICE agent calculates the priorities of host and server-reflexive candidates based on the above steps and signals these candidates in offer or answer to the remote peer. After the offer and answer are exchanged, the participating ICE agents begin pairing the candidates, ordering them into check lists to start the ICE connectivity check phase and eventually select the pair of candidates that will be used for real-time communication.

3.1. Changed Link Quality

It is possible that the characteristics of a link may change over time, and therefore the ICE agent may want to move the media to a different interface. For example, if a competing high-bandwidth flow starts or finishes its data transmission; the DSL line rate might have improved (or degraded); the link capacity may have been dynamically increased (or decreased). When link quality changes in such a fashion, the PCP Flow Extensions sends a PCP message to the endpoint. Upon receiving the message, the ICE agent may decide to move the active flow to a more suitable interface and performs ICE restart to trigger the switch over of the media streams to the new interface.

For ICE local relayed candidates, the ICE agent can switch to the more suitable interface by refreshing its allocation with the TURN server using the procedures explained in section 5 of Mobility using TURN [I-D.wing-mmusic-ice-mobility]. Thus reusing the local relayed candidate on a different interface even if the endpoint IP address changes. Therefore, the ICE agent can switch over local relayed candidate to the most suited interface that meets the requirements of the media stream. This way, even without informing the SIP server and remote peer, ICE agent can switch over a local relayed candidate to the most suited interface which meets the requested flow characteristics.

4. Multiple Interfaces

If multiple interfaces are available, the ICE agent can use PCP Flow Extensions [I-D.wing-pcp-flowdata] to determine the best path. The advantage is PCP can be used to select the most suitable interface for the media streams. When an endpoint has multiple interfaces (for example 3G, 4G, WiFi, VPN, etc.), an ICE agent can choose the interfaces for media streams according to the path characteristics, as discussed in the previous section.

4.1. Multiple Interfaces for media streams

If the requested flow characteristics for the media streams cannot be handled by a single interface but by multiple interfaces then the ICE agent performs the following steps:

- o ICE agent based on the ICE connectivity results could select multiple interfaces for the media session. For example, the ICE agent selects to send the audio stream over the WiFi access point because it offers (via PCP Flow Extensions) low delay, low packet loss and average capacity of 120 Kbps, but for the video stream it selects the 3G interface because it offers medium delay, medium packet loss and average capacity of 500Kbps.
- o Alternatively, the ICE agent on a mobile device may also want to select the best suited interface among all the available interfaces even if it does not serve the requested flow characteristics for all the media streams, so that other interfaces can be turned off to increase the battery life of cellular connected devices such as smartphones or tablets.

4.2. Availability of New Interfaces

If the available interfaces do not meet the requested flow characteristics then ICE agent can either proceed as usual using the "Recommended Formula" explained in Section 4.1.2.1 of [RFC5245] to prioritize the candidates or use the Happy Eyeballs Extension for ICE algorithm proposed in [I-D.reddy-mmusic-ice-happy-eyeballs] for dual-stack endpoint. When new interfaces become available then ICE agent can use PCP Flow Extension to find if the newly available interfaces meet the flow characteristics. When a PCP response is received from at least one of the new interfaces and if it meets the requirements, the endpoint can re-connect to the SIP proxy using the new interface. The endpoint uses the candidates indicated in the previous PCP response, it exchanges updated offer/answer to trigger ICE restart. Once the ICE processing reaches the "Completed state", the ICE endpoint can successfully switch the media session over to the new interface. The interface initially used for communication can now be

turned off without disrupting communications.

5. IANA Considerations

None.

6. Security Considerations

Security considerations discussed in [RFC6887] are to be taken into account.

7. Acknowledgements

Authors would like to thank Anca Zamfir for comments and review.

8. References

8.1. Normative References

- [I-D.wing-pcp-flowdata]
Wing, D., Penno, R., and T. Reddy, "PCP Flowdata Option", draft-wing-pcp-flowdata-00 (work in progress), July 2013.
- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.
- [RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, April 2013.

8.2. Informative References

- [I-D.penno-rtcweb-pcp]
Penno, R., Reddy, T., Wing, D., and M. Boucadair, "PCP Considerations for WebRTC Usage", draft-penno-rtcweb-pcp-00 (work in progress), May 2013.
- [I-D.reddy-mmusic-ice-happy-eyeballs]
Reddy, T., Patil, P., and D. Wing, "Happy Eyeballs Extension for ICE", draft-reddy-mmusic-ice-happy-eyeballs-03 (work in progress), October 2013.

[I-D.wing-mmusic-ice-mobility]

Wing, D., Reddy, T., Patil, P., and P. Martinsen,
"Mobility with ICE (MICE)",
draft-wing-mmusic-ice-mobility-05 (work in progress),
September 2013.

Appendix A.

A.1. Delay Factor in Discovering the Link Characteristics

Some concern has been expressed, that discovering the link characteristics may consume more time than using STUN. However, STUN will actually take more time than learning link characteristics, because a STUN request/response traverses across more routers than a PCP Flow Extension request.

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

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

Email: dwing@cisco.com

Bill VerSteeg
Cisco Systems, Inc.
5030 Sugarloaf Parkway
Lawrenceville 30044
USA

Email: billvs@cisco.com

Reinaldo Penno
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, 95134
USA

Phone:
Email: repenno@cisco.com
URI:

Varun Singh
Aalto University
School of Electrical Engineering
Otakaari 5 A
Espoo, FIN 02150
Finland

Email: varun.singh@iki.fi
URI: <http://www.netlab.tkk.fi/~varun/>