Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 14, 2008

B. Lowekamp
SIPeerior; William & Mary
D. Bryan
SIPeerior Technologies, Inc.

November 11, 2007

Using ICE to establish SIP Dialogs
draft-lowekamp-sipping-ice-for-sip-00

Status of this Memo

Copyright Notice

Abstract

This draft explores a way SIP can be extended to allow a new dialog directly between the endpoints to replace an initial dialog that had one or more proxies in the signalling path. This technique relies on ICE to perform hole punching that allows a direct connection to be used in deployments where a sip-outbound proxy or SBC is used to establish SIP connections across NAT or firewall boundaries. It can also be used to replace such a dialog with a secure connection

directly between the endpoints. This technique can be applied to traditional proxy-based SIP routing as well as to emerging P2PSIP deployments that lack centrally located proxies.

This draft describes early work that evolved from ideas initially developed for P2PSIP that are no longer being pursued. We are interested in feedback on whether there is broader interest in these techniques.

Requirements Language

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 RFC 2119 [RFC2119].

Table of Contents

1. Overview

ICE is typically used to open media streams. This draft describes how ICE can be used to open SIP signalling connections, thus "ICE for SIP." This section describes scenarios showing how an ICE for SIP extension can be used:

- Proxies can handle NAT traversal for SIP dialogs by inserting themselves into the signalling path using sip-outbound and Record-Route. ICE for SIP can be used to establish a direct connection between the endpoints after the initial setup, thus reducing the load on the proxy or proxies and the latency of the connection.

- Support from a B2BUA could allow it to remove itself from the dialog path after it is no longer interested in future call control or required for NAT traversal.

- Establishing a direct connection between the endpoints can allow for end-to-end security, which is extremely difficult to guarantee on paths where multiple proxies are involved. Here an initial connection is made using the proxies, but the dialog is replaced with a direct, secure dialog before any sensistive information is exchanged.

- Replacing a dialog originally established across a P2PSIP overlay with a direct IP connection between endpoints.

There may be other applications of this technique. In particular, if a flow established directly between endpoints can be used for future dialogs and other messages, then the proxies on the initial signalling path can be left out of those connections as well.

1.1. sip-outbound

When an inbound INVITE arrives at a proxy that supports sip-outbound[I-D.ietf-sip-outbound], the proxy is already aware that the destination UA is behind a NAT and is associated with an established flow. That edge proxy rewrites the Request-URI and forwards the message along the flow previously opened through the NAT by the client. The edge proxy adds a Record-Route header to force futher mid-dialog requests to continue to be routed through the edge proxy along the same flow.

To reduce the load on the edge proxy, ICE for SIP allows the two endpoints (or other proxies along the path) to establish a direct connection for further mid-dialog requests. When a UA sends a request to open a dialog, it includes an ice-sip tag in its Via. Similarly, the proxy adds the same tag to its Record-Route header.

After the initial dialog is established, the answering endpoints inspects the components of the path. ICE for SIP may be used to replace the initial dialog if the initial endpoint added an ice-sip tag to its Via header and each proxy along the path that inserted a Record-Route header indicated its support for ice-sip through a tag in its Record-Route URI. Note that this requirement is to ensure that policy enforced by intermediate proxies is not bypassed by the replacement dialog rather than by a technical requirement that the intermediate proxies must meet. A proxy that supports ICE for SIP but is unwilling to allow the dialog to be replaced with a direct path would not insert an ice-sip tag into its Record-Route header for that particular dialog.

Open issue: should it be possible for intermediate proxies to make use of this feature to remove other intermediate proxies from the path even if the endpoints do not themselves support ice-sip, i.e. shorten the path even if a direct connection is not possible?

To replace the initial dialog, the answering endpoint initiates a re-INVITE with an ICE SDP that specifies a media type of control/sip. The endpoints then perform ICE negotiation and, if successful, the offerer sends an INVITE across the newly established end-to-end flow with a Replaces header that indicates the original dialog is being replaced[RFC3891].

Open issue: what about dialogs not established by an INVITE?

Open issue: Could a flow established be use for future dialogs or non-dialog use such as MESSAGE? Should it be possible to specify an INVITE that specifically requests this behavior so that on-path proxies can process/reject it if they want to be aware of future dialogs? Technically this is rather simple once the direct flow is open, it's just a question of whether it might violate a proxy's policy requirements. Perhaps in addition to ice-sip there should be another tag or the tag should have an option to indicate whether only this or future dialogs may be directly routed?

1.2. B2BUA

In an SBC type deployment the endpoints are typically not aware that there is a way the path could be optimized because they do not see end-to-end headers. However, if the B2BUA indicates its support for ice-sip as above, and all other elements on the path support ice-sip, that B2BUA may initiate dialog replacement even if it appears to the other endpoint that there are no other elements that inserted themselves into the path with Record-Route headers.

Replacing the new dialog is conceptually simple, except that the

existing dialog presumably has a different dialog id (call-id, to-tag, and from-tag) on either side of the B2BUA. Therefore, a direct end-to-end INVITE with a Replaces header would not work. Instead, a REFER has more appropriate semantics and could be used instead.

Open issue: it is unclear whether it would be worth specifying such behavior for a B2BUA acting as an SBC because it might make more sense for such a device to be redeployed as a sip-outbound capable device that could more naturally implement ICE for SIP and not worry about the complexity of addressing this situation. In particular, if an SBC is used to provide demarcation and intended to hide the internal network, rather than just facilitating NAT traversal, a direct connection would not be appropriate.

Open issue: this technique could be used to bypass a Controller in 3pcc call flows. Is there interest in such a capability?

Open issue: Rather than using REFER, it might be better to provide a technique where UAs implementing the ice-sip extension identify that there is a B2BUA involved in the initial re-INVITE and rely on ICE's authentication from the SDP in the re-INVITE to connect the old and new dialogs.

1.3. Secure Connection

Ensuring the security of an end-to-end SIP dialog in the presence of multiple proxies is a difficult challenge, and there is no way a UA can be certain that a message was delivered securely along each hop [I-D.ietf-sip-sips]. In this case, the techniques of 1.1 can be used to ensure security by establishing a direct TLS or DTLS connection between the endpoints. Rather than establishing an initial dialog with an INVITE specifying media to be exchanged, the initial INVITE can merely specify a control/sip media type, initiating the creation of a direct, secure dialog that can be used for future exchanges and real media.

1.4. P2PSIP

P2PSIP, by definition, relies on end-to-end connections between its peers for SIP dialogs. Multiple mechanisms have been proposed for establishing these dialogs, with some proposals suggesting multiple methods [I-D.bryan-p2psip-reload][I-D.matthews-p2psip-hip-hop]:

1. Direct connection between peers, assuming that all peers will accept direct incoming connections.

2. Indirect connection established through an intermediary, typically using ICE. The intermediary could either be a single

entity, if one with appropriate connectivity can be located, or the P2P overlay network itself.

3. Tunneled connection relying on the overlay for transport of SIP datagrams.

4. HIP-HOP relies on an entirely different technique of using the connectivity obtained by using HIP to route SIP messages.

The first technique is obviously of limited applicability in scenarios that range beyond a single LAN. The second technique works well, but imposes the ICE setup delay on the new connection before the actual SIP message can be sent. The third technique avoids the initial ICE setup delay, but establishes the dialog across the overlay, resulting in the overlay's routing latency being added to each message exchanged in the SIP dialog.

The tradeoff between the second and third technique is that the first trades initial delay for a direct dialog connection, whereas the third has lower initial delay, but an indirect connection for the entire dialog. Although the second technique relies on the use of ICE to establish a SIP dialog, it does not require use of the specification in this draft because it concerns only establishing a new dialog and is expected to be encoded in a custom representation, rather than SDP.

The third technique's shortcoming of higher per-message latency can be resolved by applying ICE for SIP to replace the initial overlay-routed dialog with a direct dialog. Thus, the initial dialog can be established quickly by routing across the overlay and deferring ICE negotiation until the dialog is established. If ICE negotiation goes slowly or fails, the overlay-routed dialog can continue to be used. Otherwise, it will be replaced by the end-to-end dialog.

2. Extensions to SIP

The initial requester SHOULD include an ice-sip tag in their via to indicate a willingness to accept ICE negotiation for a replacement dialog.

Any proxy that inserts a Record-Route for itself SHOULD add ice-sip tag to its URI in the Record-Route header if it wishes to allow the dialog to be replaced with a direct dialog that bypasses itself. If a proxy wishes to be involved in all future messages in the dialog, it MUST NOT include an ice-sip tag in its Record-Route header.

The answerer MUST NOT initiate a request for a replacement dialog

unless the initial Via and all Record-Route URIs contain an ice-sip tag.

3. ICE Negotiation

ICE negotiation is handled as described in ICE [I-D.ietf-mmusic-ice] and ICE-TCP [I-D.ietf-mmusic-ice-tcp]. ICE for SIP uses SDP to encode its ICE offers and answers because all SIP implementations already implement SDP and those implementing ICE will support encoding ICE offers in SDP. The following changes are made for ICE for SIP negotiations from ICE for media:

- Timers will be set as specified in Section 16.2 of ICE [I-D.ietf-mmusic-ice]for non-RTP sessions.

- The SDP's "m=" line will specify the media type as "control" and the media format as "sip". The transport field will be either "tcp" or "udp". SDP [RFC4566]

- Specification of encryption requirements in the SDP is an open issue being addressed in the MMUSIC working group. We will use those techniques when they are finalized.

- The SDP MUST NOT include an "a=recvonly", "a=sendonly", "a=inactive", or a "0.0.0.0" specification.

- A relay candidate SHOULD NOT be included in the SDP. As the dialog has an existing path through proxies, there should be no reason to switch to a different method of relaying.

If ICE negotiation fails, then the re-INVITE has failed and the UAs will continue to use the existing dialog. The UAs MUST NOT attempt to use a default destination.

Open Issue: should a default destination of 0.0.0.0/0 be specified?

Open Issue: dsip-nat-traversal [I-D.matthews-p2psip-dsip-nat-traversal] specified media type application/sip, but this seems inappropriate as it doesn't meet the definition of "application" data that is to be presented to a user from SDP [RFC2327]. The former media type of "control" seems to be more appropriate for a SIP signalling connection.

4. IANA Considerations

TBD

5. Security Considerations

The technique described in this draft poses policy issues in that it allows SIP UAs to bypass proxies that would ordinarily be in the path between those UAs. However, because the dialog will not be replaced unless each proxy in the path that would be kept in the dialog authorizes such a change by inserting an "ice-sip" tag, policy requirements to keep a proxy in the path are maintained.

Attacks on the ICE negotiation are addressed in ICE [I-D.ietf-mmusic-ice]. ICE is best secured by securing the initial SIP dialog, which secures the initial SDP exchange.

The replacement dialog should also be secured as a sips connection with TLS or DTLS. Because the endpoints have been authenticated with ICE, Diffie-Hellman can be used or possibly TLS-PSK could be used with the ice-pwd values from the SDP used to form the key.

There are likely other security risks that are have not yet been considered.

6. Acknowledgements

The idea for using INVITE to establish a new SIP session originated in the earliest work on P2PSIP [I-D.bryan-sipping-p2p] as a technique for establishing connections between peers in the overlay. Further work refined the concept for NAT traversal for a P2PSIP overlay [I-D.matthews-p2psip-dsip-nat-traversal][I-D.matthews-p2psip-bootstrap-mechanisms]. Jonathan Rosenberg pointed out that the technique might have applications for regular SIP deployments in addition to P2PSIP. Thanks to Alan Johnston and special thanks to Philip Matthews for many conversations on NAT traversal for P2PSIP.

7. References

7.1. Normative References

[I-D.ietf-mmusic-ice]
Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", draft-ietf-mmusic-ice-19 (work in progress), October 2007.

[I-D.ietf-mmusic-ice-tcp]
Rosenberg, J., "TCP Candidates with Interactive Connectivity Establishment (ICE",

draft-ietf-mmusic-ice-tcp-04 (work in progress), July 2007.

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

[RFC3891] Mahy, R., Biggs, B., and R. Dean, "The Session Initiation Protocol (SIP) "Replaces" Header", RFC 3891, September 2004.

[RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.

7.2. Informative References

[I-D.bryan-p2psip-reload]
Bryan, D., "REsource LOcation And Discovery (RELOAD)", draft-bryan-p2psip-reload-01 (work in progress), July 2007.

[I-D.bryan-sipping-p2p]
Bryan, D., "A P2P Approach to SIP Registration and Resource Location", draft-bryan-sipping-p2p-03 (work in progress), October 2006.

[I-D.ietf-sip-outbound]
Jennings, C. and R. Mahy, "Managing Client Initiated Connections in the Session Initiation Protocol (SIP)", draft-ietf-sip-outbound-10 (work in progress), July 2007.

[I-D.ietf-sip-sips]
Audet, F., "The use of the SIPS URI Scheme in the Session Initiation Protocol (SIP)", draft-ietf-sip-sips-06 (work in progress), August 2007.

[I-D.matthews-p2psip-bootstrap-mechanisms]
Cooper, E., "Bootstrap Mechanisms for P2PSIP", draft-matthews-p2psip-bootstrap-mechanisms-00 (work in progress), February 2007.

[I-D.matthews-p2psip-dsip-nat-traversal]
Cooper, E., "NAT Traversal for dSIP", draft-matthews-p2psip-dsip-nat-traversal-00 (work in progress), February 2007.

[I-D.matthews-p2psip-hip-hop]
Cooper, E., "A Distributed Transport Function in P2PSIP using HIP for Multi-Hop Overlay Routing",

draft-matthews-p2psip-hip-hop-00 (work in progress), June 2007.

[RFC2327] Handley, M. and V. Jacobson, "SDP: Session Description Protocol", RFC 2327, April 1998.

Authors' Addresses

Bruce B. Lowekamp
SIPeerior; William & Mary
3000 Easter Circle
Williamsburg, VA 23188
USA

Phone: +1 757 565 0101
Email: lowekamp@sipeerior.com

David A. Bryan
SIPeerior Technologies, Inc.
3000 Easter Circle
Williamsburg, VA 23188
USA

Phone: +1 757 565 0101
Email: dbryan@sipeerior.com

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).