MMUSIC
Internet-Draft
Expires: September 30, 2006

J. Rosenberg
Cisco Systems

March 29, 2006

Interactive Connectivity Establishment (ICE): A Methodology for Network Address Translator (NAT) Traversal for Offer/Answer Protocols

draft-ietf-mmusic-ice-08

Status of this Memo

Copyright Notice

Abstract

This document describes a protocol for Network Address Translator (NAT) traversal for multimedia session signaling protocols based on the offer/answer model, such as the Session Initiation Protocol (SIP). This protocol is called Interactive Connectivity Establishment (ICE). ICE makes use of the Simple Traversal of UDP through NAT (STUN), applying its binding discovery, connectivity check and relay usages.

Table of Contents

1. Introduction

RFC 3264 [4] defines a two-phase exchange of Session Descrption Protocol (SDP) messages [5] for the purposes of establishment of multimedia sessions. This offer/answer mechanism is used by protocols such as the Session Initiation Protocol (SIP) [2].

Protocols using offer/answer are difficult to operate through Network Address Translators (NAT). Because their purpose is to establish a flow of media packets, they tend to carry IP addresses within their messages, which is known to be problematic through NAT [17]. The protocols also seek to create a media flow directly between participants, so that there is no application layer intermediary between them. This is done to reduce media latency, decrease packet loss, and reduce the operational costs of deploying the application. However, this is difficult to accomplish through NAT. A full treatment of the reasons for this is beyond the scope of this specification.

Numerous solutions have been proposed for allowing these protocols to operate through NAT. These include Application Layer Gateways (ALGs), the Middlebox Control Protocol [19], Simple Traversal of UDP through NAT (STUN) [16] and its revision [13], the STUN Relay Usage [14], and Realm Specific IP [20] [21] along with session description extensions needed to make them work, such as the Session Description Protocol (SDP) [5] attribute for the Real Time Control Protocol (RTCP) [1]. Unfortunately, these techniques all have pros and cons which make each one optimal in some network topologies, but a poor choice in others. The result is that administrators and implementors are making assumptions about the topologies of the networks in which their solutions will be deployed. This introduces complexity and brittleness into the system. What is needed is a single solution which is flexible enough to work well in all situations.

This specification provides that solution for media streams established by signaling protocols based on the offer-answer model. It is called Interactive Connectivity Establishment, or ICE. ICE makes use of STUN and its relay extension, commonly called TURN, but uses them in a specific methodology which avoids many of the pitfalls of using any one alone.

2. Terminology

Several new terms are introduced in this specification:

Agent: As defined in RFC 3264, an agent is the protocol implementation involved in the offer/answer exchange. There are two agents involved in an offer/answer exchange.

Peer: From the perspective of one of the agents in a session, its peer is the other agent. Specifically, from the perspective of the offerer, the peer is the answerer. From the perspective of the answerer, the peer is the offerer.

Transport Address: The combination of an IP address and port.

Local Transport Address: A local transport address is a transport address that has been allocated from the operating system on the host. This includes transport addresses obtained through Virtual Private Networks (VPNs) and transport addresses obtained through Realm Specific IP (RSIP) [20] (which lives at the operating system level). Transport addresses are typically obtained by binding to an interface.

m/c line: The media and connection lines in the SDP, which together hold the transport address used for the receipt of media.

Derived Transport Address: A derived transport address is a transport address which is derived from a local transport address. The derived transport address is related to the associated local transport address in that packets sent to the derived transport address are received on the socket bound to its associated local transport address. Derived addresses are obtained using protocols like STUN, and more generally, any UNSAF protocol [22].

Reflexive Transport Address: As defined in [13], a transport address learned by a client which identifies that client as seen by another host on an IP network, typically a STUN server. When there is an intervening NAT between the client and the other host, the reflexive transport address represents the binding allocated to the client on the public side of the NAT. Reflexive transport addresses are learned from the XOR-MAPPED-ADDRESS attribute in STUN Binding Responses and Allocate Responses [14], and are a type of derived transport address.

Server Reflexive Transport Address: A server reflexive transport address is a reflexive address that is reflected off of a server, distinct from the peer, whose address is configured or learned by the client prior to an offer/answer exchange.

Peer Reflexive Transport Address: A peer reflexive transport address is a reflexive address that is reflected off of the peer. Peer reflexive transport addresses are learned by connectivity checks.

Relayed Transport Address: A transport address that terminates on a server, and is forwarded towards the client. The STUN Allocate Request can be used to obtain a relayed transport address, for example.

Associated Local Transport Address: When a peer sends a packet to a transport address, the associated local transport address is the local transport address at which those packets will actually arrive. For a local transport address, its associated local transport address is the same as the local transport address itself. For reflexive and relayed transport addresses, however, they are not the same. The associated local transport address is the one from which the reflexive or relayed transport was derived.

Candidate: A sequence of transport addresses that form an atomic set for usage with a particular media session. Here, atomic means that all of transport addresses in the candidate need to work before the candidate will be used for actual media transport. In the case of RTP, there can be one or more transport addresses per candidate. In the most common case, there are two - one for RTP, and another for RTCP. If the agent doesn't use RTCP, there would be just one. If Generic Forward Error Correction (FEC) [18] is in use, there may be more than two. The transport addresses that compose a candidate are all of the same type - local, server reflexive, peer reflexive or relayed.

Local Candidate: A candidate whose transport addresses are local transport addresses.

Server Reflexive Candidate: A candidate whose transport addresses are server reflexive transport addresses.

Peer Reflexive Candidate: A candidate whose transport addresses are peer reflexive transport addresses.

Relayed Candidate: A candidate whose transport addresses are relayed transport addresses.

Generating Candidate: The candidate from which a peer reflexive candidate is derived.

Active Candidate: The candidate that is in use for exchange of media. This is the one that an agent places in the m/c line of an offer or answer.

Candidate ID: An identifier for a candidate.

Component: When a media stream, and as a consequence, its candidate, require several IP addresses and ports to work atomically, each of the constituent IP addresses and ports represents a component of that media stream. For example, RTP-based media streams typically have two components - one for RTP, and one for RTCP.

Component ID: An integer, starting with one within each candidate and incrementing by one for each component, which identifies the component.

Transport Address ID (tid): An identifier for a transport address, formed by concatenating the candidate ID with the component ID, separated by a "colon".

Candidate Pair: The combination of a candidate from one agent along with a candidate from its peer.

Native Candidate: From the perspective of each agent, the candidate in a candidate pair which represents a set of addresses obtained by that agent.

Remote Candidate: From the perspective of each agent, the candidate in a candidate pair which represents the set of addresses obtained by that agents peer.

Transport Address Pair: The combination of the transport address for one component of a candidate with the transport address of the same component for the matching candidate in a candidate pair.

Transport Address Pair ID: An identifier for a transport address pair. Formed by concatenating the native transport address ID with the remote transport address ID, separated by a "colon".

Matching Transport Address Pair: When a STUN Binding Request is received on a local transport address, the matching transport address pair is the transport address pair whose connectivity is being checked by that Binding Request.

Candidate Pair Priority Ordering: An ordering of candidate pairs based on a combination of the qvalues of each candidate and the candidate IDs of each candidate.

Candidate Pair Check Ordering: An ordering of candidate pairs that is similar to the candidate pair priority ordering, except that the active candidate appears at the top of the list, regardless of its priority.

Transport Address Pair Check Ordering: An ordering of transport address pairs that determines the sequence of connectivity checks performed for the pairs.

Transport Address Pair Count: The number of transport address pairs in a candidate pair. This is equal to the minimum of the number of transport addresses in the native candidate and the number of transport addresses in the remote candidate.

3. Overview of ICE

ICE makes the fundamental assumption that clients exist in a network of segmented connectivity. This segmentation is the result of a number of addressing realms in which a client can simultaneously be connected. We use "realms" here in the broadest sense. A realm is defined purely by connectivity. Two clients are in the same realm if, when they exchange the addresses each has in that realm, they are able to send packets to each other. This includes IPv6 and IPv4 realms, which actually use different address spaces, in addition to private networks connected to the public Internet through NAT.

The key assumption in ICE is that a client cannot know, apriori, which address realms it shares with any peer it may wish to communicate with. Therefore, in order to communicate, it has to try connecting to addresses in all of the realms.

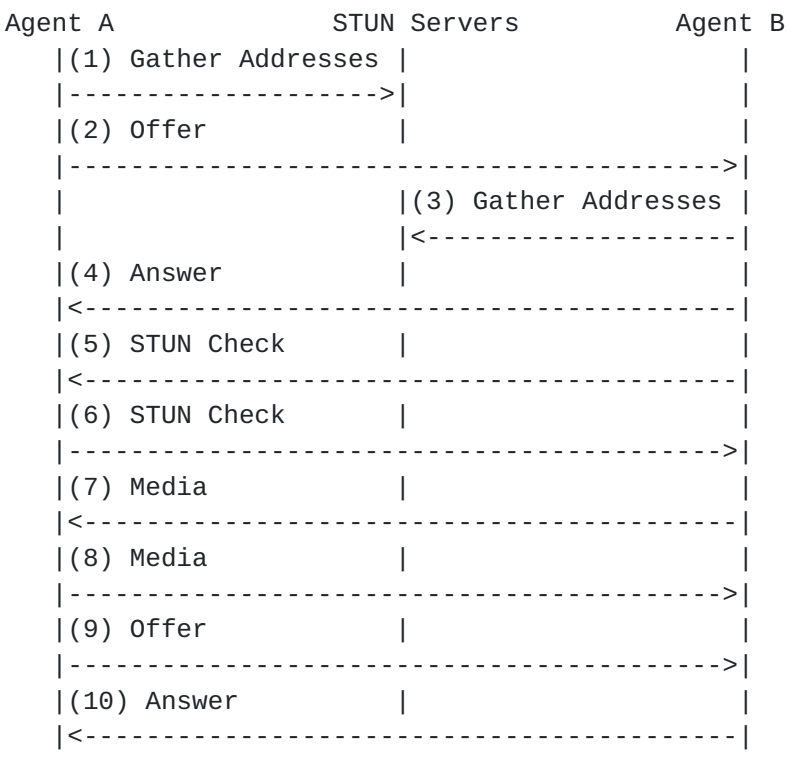

Figure 1

The basic flow of operation for ICE is shown in Figure 1. Before the offerer establishes a session, it obtains local transport addresses from its operating system on as many interfaces as it has access to. These interfaces can include IPv4 and IPv6 interfaces, in addition to Virtual Private Network (VPN) interfaces or ones associated with RSIP. It then obtains transport addresses for the media from each interface. Though ICE can support any type of transport protocol, this specification only defines mechanisms for UDP. In addition, the agent obtains server reflexive and relayed transport addresses. These are usually obtained through a single STUN Allocate request, which provides both. These requests are paced at a fixed rate in order to limit network load and avoid NAT overload. The local, server reflexive and relayed transport addresses are formed into candidates, each of which represents a possible set of transport addresses that might be viable for a media stream.

Each candidate is listed in a set of a=candidate attributes in the offer. Each candidate is given a priority. Priority is a matter of local policy, but typically, lowest priority would be given to relayed transport addresses. Each candidate is also assigned a distinct ID, called a candidate ID.

The agent will choose one of its candidates as its active candidate

for inclusion in the connection and media lines in the offer. Media can be sent to this candidate immediately following its validation. Media can also be sent to a candidate that is not active but has been validated. Media is not sent without validation in order to avoid denial-of-service attacks. In particular, without ICE, an offerer can send an offer to another agent, and list the IP address and port of a target in the offer. If the agent is an automata that answers a call automatically, it will do so and then proceed to send media to the target. This provides substantial packet amplifications. ICE fixes this by requiring that an agent never send media packets unless it has sent a STUN message towards the target of the RTP packets, and received a reply from that target Section 7.13.

The offer is then sent to the answerer. This specification does not address the issue of how the signaling messages themselves traverse NAT. It is assumed that signaling protocol specific mechanisms are used for that purpose. The answerer follows a similar process as the offerer followed; it obtains addresses from local interfaces, obtains derived transport addresses from those, and then groups them into candidates for inclusion in a=candidate attributes in the answer. It picks one candidate as its active candidate and places it into the m/c line in the answer.

Once the offer/answer exchange has completed, both agents pair up the candidates, and then determine an ordered set of transport address pairs. This ordering is based primarily on the priority of the candidates, with the exception of the active candidate, whose addresses are at the top of the list. Both agents start at the top of this list, beginning a connectivity check for that transport address pair. At a fixed interval, checks for the next transport address on the list begin. This results in a pacing of the connectivity checks. These connectivity checks are performed through peer-to-peer STUN requests, sent from one agent to the other. In addition to pacing the checks out at regular intervals, the offerer will generate a connectivity check for a transport address pair when it receives one from its peer. As soon as the active candidate has been verified by the STUN checks, media can begin to flow. Once a higher priority candidate has been verified by the offerer, it ceases additional connectivity checks, begins using that candidate for media, and sends an updated offer which promotes this higher priority candidate to the m/c-line. That candidate is also listed in a=candidate attributes, resulting in periodic STUN keepalives through the duration of the media session.

If an agent receives a STUN connectivity check with a new source IP address and port, or a response to such a check with a new reflexive transport address (obtained from the XOR-MAPPED-ADDRESS attribute), this new address might be a viable candidate for the receipt of

media. This happens when there is a NAT with an address dependent or address and port dependent mapping property [37] between the agents. In such a case, the agents algorithmically construct a new candidate. Like other candidates, connectivity checks begin for it, and if they succeed, its transport addresses can be used for receipt of media by promoting it to the m/c-line.

The gathering of addresses and connectivity checks take time. As a consequence, in order to have minimal impact on the call setup time or post-pickup delay for SIP, these offer/answer exchanges and checks happen while the call is ringing.

4. Sending the Initial Offer

When an agent wishes to begin a session by sending an initial offer, it starts by gathering transport addresses, as described in Section 7.1. This will produce a set of candidates, including local ones, server reflexive ones, and relayed ones.

This process of gathering candidates can actually happen at any time before sending the initial offer. A agent can pre-gather transport addresses, using a user interface cue (such as picking up the phone, or entry into an address book) as a hint that communications is imminent. Doing so eliminates any additional perceivable call setup delays due to address gathering.

When it comes time to offer communications, the agent determines a priority for each candidate and identifies the active candidate that will be used for receipt of media, as described in Section 7.2.

The next step is to construct the offer message. For each media stream, it places its candidates into a=candidate attributes in the offer and puts its active candidate into the m/c line. The process for doing this is described in Section 7.3. The offer is then sent.

5. Receipt of the Offer and Generation of the Answer

Upon receipt of the offer message, the agent checks if the offer contains any a=candidate attributes. If the offer does, the offerer supports ICE. In that case, it starts gathering candidates, as described in Section 7.1, and prioritizes them as described in Section 7.2. This processing is done immediately on receipt of the offer, to prepare for the case where the user should accept the call, or early media needs to be generated. By gathering candidates (and performing connectivity checks) while the user is being alerted to the request for communications, session establishment delays are reduced.

The agent then constructs its answer, encoding its candidates into a=candidate attributes and including the active one in the m/c-line, as described in Section 7.3. The agent then forms candidate pairs as described in Section 7.4. These are ordered as described in Section 7.5. The agent then begins connectivity checks, as described in Section 7.6. It follows the logic in Section 7.10 on receipt of Binding Requests and responses to learn new candidates from the checks themselves.

Transmission of media is performed according to the procedures in Section 7.13.

6. Processing the Answer

There are two possible cases for processing of the answer. If the answerer did not support ICE, the answer will not contain any a=candidate attributes. As a result, the offerer knows that it cannot perform its connectivity checks. In this case, it proceeds with normal media processing as if ICE was not in use. However, it SHOULD send media with the symmetric property described in Section 7.13, and follow the keepalive procedures in Section 7.12.

If the answer contains candidates, it implies that the answerer supports ICE. The offerer then forms candidate pairs as described in Section 7.4. These are ordered as described in Section 7.5. The agent then begins connectivity checks, as described in Section 7.6. It follows the logic in Section 7.10 on receipt of Binding Requests and responses to learn new candidates from the checks themselves.

Transmission of media is performed according to the procedures in Section 7.13.

7. Common Procedures

This section discusses procedures that are common between offerer and answerer.

7.1 Gathering Candidates

An agent gathers candidates when it believes that communications is imminent. For offerers, this occurs before sending an offer (Section 4). For answerers, it occurs before sending an answer (Section 5).

Each candidate has one or more components, each of which is associated with a sequence number, starting at 1 for the first component of each candidate, and incrementing by 1 for each additional component within that candidate. These components

represent a set of transport addresses for which connectivity must be validated. For a particular media stream, all of the candidates SHOULD have the same number of components. The number of components that are needed are a function of the type of media stream. All of the components in a candidate MUST be of the same type - server reflexive, relayed, or local, and obtained from the same server in the case of server reflexive or relayed candidates. For local candidates, each component MUST be obtained from the same interface.

For traditional RTP-based media streams, it is RECOMMENDED that there be two components per candidate - one for RTP and one for RTCP. The component with the component ID of 1 MUST be RTP, and the one with component ID of 2 MUST be RTCP. If an agent doesn't implement RTCP, it SHOULD have a single component for the RTP stream (which will have a component ID of 1 by definition). Each component of a candidate has a single transport address.

The first step is to gather local candidates. Local candidates are obtained by binding to ephemeral ports on an interface (physical or virtual, including VPN interfaces) on the host. The process for gathering local candidates depends on the transport protocol. Procedures are specified here for UDP. Extensions to ICE that define procedures for other transport protocols MUST specify how local transport addresses are gathered.

For each UDP media stream the agent wishes to use, the agent SHOULD obtain a set of candidates (one for each interface) by binding to N ephemeral UDP ports on each interface, where N is the number of components needed for the candidate. For RTP, N is typically two. If a host has K local interfaces, this will result in K candidates for each UDP stream, requiring K*N local transport addresses.

Once the agent has obtained local candidates, it obtains candidates with derived transport addresses. The process for gathering derived candidates depends on the transport protocol. Procedures are specified here for UDP. Extensions to ICE that define procedures for other transport protocols MUST specify how derived transport addresses are gathered.

Agents which serve end users directly, such as softphones, hardphones, terminal adapters and so on, MUST implement the STUN Binding Discovery usage and SHOULD use it to obtain server reflexive candidates. These devices SHOULD implement the STUN Relay usage, and SHOULD use its Allocate request to obtain both server reflexive and relayed candidates. They MAY implement and MAY use other protocols that provide server reflexive or relayed transport addresses, such as TEREDO [33].

The requirement to use the relay Usage is at SHOULD strength to allow for provider variation. If it is not to be used, it is RECOMMENDED that it be implemented and just disabled through configuration, so that it can re-enabled through configuration if conditions change in the future.

Agents which represent network servers under the control of a service provider, such as gateways to the telephone network, media servers, or conferencing servers that are targeted at deployment only in networks with public IP addresses MAY use the STUN Binding Discovery usage and relay usage, or other similar protocols to obtain candidates.

> Why would these types of endpoints even bother to implement ICE? The answer is that such an implementation greatly facilitates NAT traversal for clients that connect to it. The ability to process STUN connectivity checks allows for clients to obtain peer reflexive transport addresses that can be used by the network server to reach them without a relay, even through NATs with restrictive mapping and filtering policies. Furthermore, implementation of the STUN connectivity checks allows for NAT bindings along the way to be kept open. ICE also provides numerous security properties that are independent of NAT traversal, and would benefit any multimedia endpoint. See Section 13 for a discussion on these benefits.

Obtaining derived candidates requires transmission of packets which have the effect of creating bindings on NAT devices between the client and the STUN servers. Experience has shown that many NAT devices have upper limits on the rate at which they will create new bindings. Furthermore, transmission of these packets on the network makes use of bandwidth and needs to be rate limited by the agent. As a consequence, a client SHOULD pace its STUN transactions, such that the start of each new transaction occurs at least Ta seconds after the start of the previous transaction. The value of Ta SHOULD be configurable, and SHOULD have a default of 50ms. Note that this pacing applies only to the start of a new transaction; pacing of retransmissions within a STUN transaction is governed by the retransmission rules defined by STUN.

Derived candidates can be obtained from the STUN Binding Discovery usage or the STUN Relay usage. The latter is preferred since it will provide the client with both a server reflexive and a relayed transport address with a single transaction. It is possible that some STUN servers will only support the Relay usage or only the Binding Discovery usage, in which case a client might be configured with different servers depending on the usage.

To obtain both server reflexive and relayed candidates using the STUN Relay Usage, the client takes a local UDP candidate, and for each configured STUN server, produces both candidates. It is anticipated that clients may have a multiplicity of STUN servers configured or discovered in network environments where there are multiple layers of NAT, and that layering is known to the provider of the client. To obtain these candidates, for each configured STUN server, the client initiates an Allocate Request transaction using the procedures of Section 8.1.2 of [14] from each transport address of a particular local candidate. The Allocate Response will provide the client with its server reflexive transport address (obtained from the XOR-MAPPED-ADDRESS attribute) and its relayed transport address in the RELAY-ADDRESS attribute. Once the Allocate requests have given a client a relayed transport address for all transport addresses in a relayed candidate, there is no reason for a client to obtain further relayed candidates through the same STUN server. Thus, if there are other local candidates from which the client has not yet obtained relayed transport address, the client SHOULD NOT bother to obtain them. Instead, it SHOULD use the STUN Binding Discovery usage and obtain just server reflexive addresses from that STUN server. The order in which local candidates are tried against the STUN server to obtain relayed candidates is a matter of local policy.

To obtain server reflexive candidates using the STUN Binding Discovery usage, the client takes a local UDP candidate, and for each configured STUN server, produces a server reflexive candidate. To produce the server reflexive candidate from the local candidate, it follows the procedures of Section 12.2 of [13] for each local transport address in the local candidate. The Binding Response will provide the client with its server reflexive transport address. If the client had K local candidates, this will produce S*K server reflexive candidates, where S is the number of STUN servers.

Since a client will pace its STUN transactions (both Binding and Allocate requests) at a total rate of one new transaction every Ta seconds, it will take a certain amount of time to complete the address gathering phase. It is RECOMMENDED that implementations have a configurable upper bound on the total amount of time allotted to address gathering. Any transactions not completed at that point SHOULD be abandoned, but MAY continue and be used in an updated offer once they complete. A default value of 5s is RECOMMENDED. Since the total number of allocations that could be done (based on the number of STUN servers and local interfaces) might exceed this value, clients SHOULD prioritize their local candidates and STUN servers, performing transactions from the highest priority local candidates to the highest priority STUN servers first. A STUN server would typically be higher priority if it supports the STUN Relay Usage, since such a server provides two transport addresses with one

transaction.

Once the allocations are complete, any redundant candidates are discarded. Candidate A is redundant with candidate B if the transport addresses for each component of each component match, and each component of their associated local candidates match. For example, consider a set of candidates with a single component. One candidate is a local candidate, and its one component has a transport address of 10.0.1.1:4458. A reflexive transport address is derived from this local transport address, producing a 10.0.1.1:4458. These two candidates are identical, and also have identical associated local transport addresses, so they are redundant.

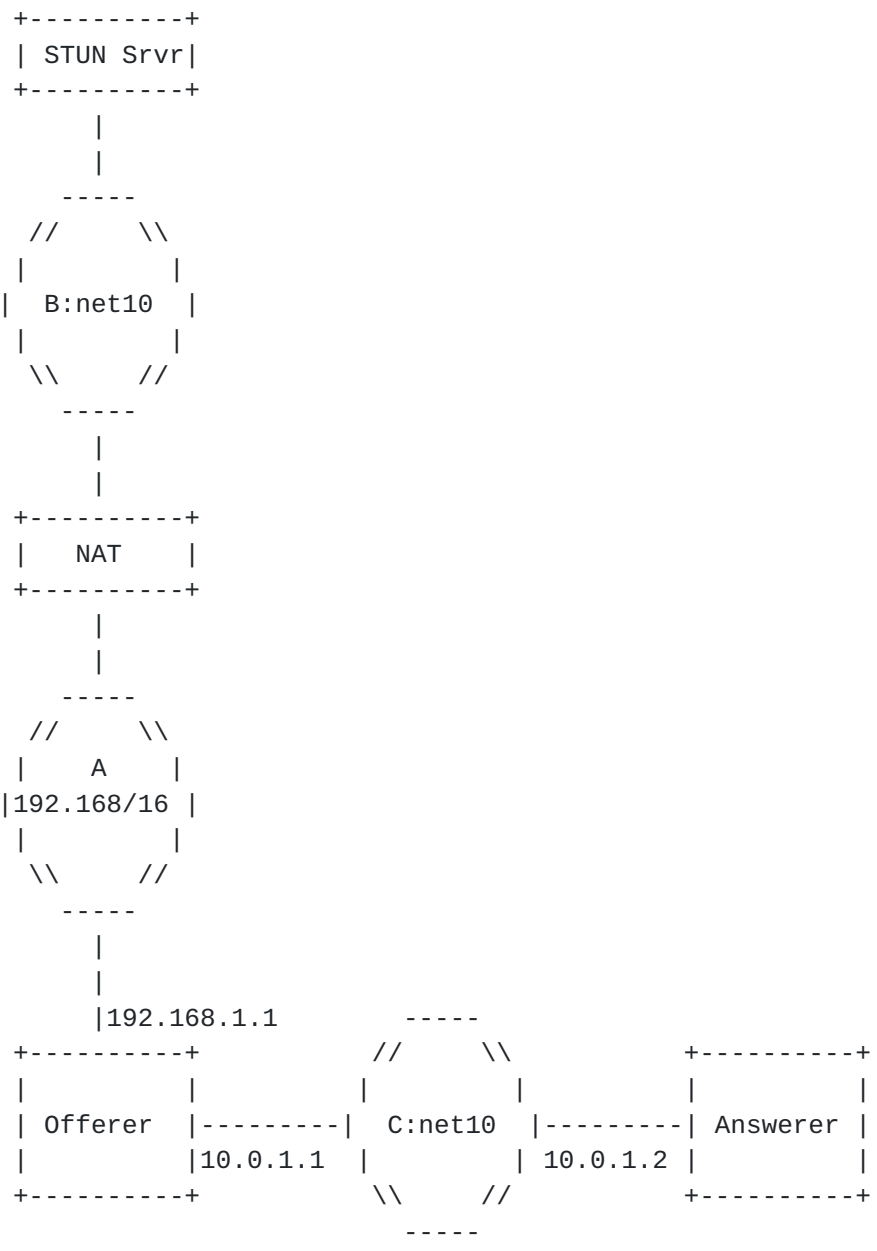

Figure 2

Consider the more complicated case of Figure 2. In this case, the offerer is multi-homed. It has one interface, 10.0.1.1, on network C, which is a net 10 private network. The Answerer is on this same network. The offerer is also connected to network A, which is 192.168/16. The offerer has an interface of 192.168.1.1 on this network. There is a NAT on this network, natting into network B, which is another net10 private network, but not connected to network C. There is a STUN server on network B.

The offerer obtains local transport address on its interface on network C (10.0.1.1:2498) and a local transport address on its interface on network A (192.168.1.1:3344). It performs a STUN query to its configured STUN server from 192.168.1.1:3344. This query passes through the NAT, which happens to assign the binding 10.0.1.1:2498. The STUN server reflects this in the STUN Binding Response. Now, the offerer has obtained a candidate with a transport address it already has (10.0.1.1:2498), but from a new interface. It therefore keeps it. When it performs its connectivity checks, the offerer will end up sending packets from both interfaces, and those sent from its interface on network C will succeed.

7.2 Prioritizing the Candidates and Choosing an Active One

The prioritization process takes the set of candidates and associates each with a priority. This priority reflects the desire that the agent has to receive media at that candidate, and is assigned as a value from 0 to 1 (1 being most preferred). Priorities are ordinal, so that their significance is only meaningful relative to other candidates from that agent for a particular media stream. Candidates MAY have the same priority. However, it is RECOMMENDED that each candidate have a distinct priority. Doing so improves the efficiency of ICE.

This specification makes no normative statements on how the prioritization is done. However, some useful guidelines are suggested on how such a prioritization can be determined.

One criteria for choosing one candidate over another is whether or not that candidate involves the use of an intermediary. That is, if media is sent to that candidate, will the media first transit an intermediate server before being received. Relayed candidates are clearly one type of candidates that involve an intermediary. Another are local candidates associated with a VPN server. When media is transited through an intermediary, it can increase the latency between transmission and reception. It can increase the packet losses, because of the additional router hops that may be taken. It may increase the cost of providing service, since media will be routed in and right back out of an intermediary run by the provider. If these concerns are important, candidates with this property can be listed with lower priority.

Another criteria for choosing one candidate over another is IP address family. ICE works with both IPv4 and IPv6. It therefore provides a transition mechanism that allows dual-stack hosts to prefer connectivity over IPv6, but to fall back to IPv4 in case the v6 networks are disconnected (due, for example, to a failure in a 6to4 relay) [25]. It can also help with hosts that have both a

native IPv6 address and a 6to4 address. In such a case, higher priority could be afforded to the native v6 address, followed by the 6to4 address, followed by a native v4 address. This allows a site to obtain and begin using native v6 addresses immediately, yet still fallback to 6to4 addresses when communicating with agents in other sites that do not yet have native v6 connectivity.

Another criteria for choosing one candidate over another is security. If a user is a telecommuter, and therefore connected to their corporate network and a local home network, they may prefer their voice traffic to be routed over the VPN in order to keep it on the corporate network when communicating within the enterprise, but use the local network when communicating with users outside of the enterprise.

Another criteria for choosing one address over another is topological awareness. This is most useful for candidates that make use of relays. In those cases, if an agent has preconfigured or dynamically discovered knowledge of the topological proximity of the relays to itself, it can use that to select closer relays with higher priority.

There may be transport-specific reasons for preferring one candidate over another. In such a case, specifications defining usage of ICE with other transport protocols SHOULD document such considerations.

Once the candidates have been prioritized, one may be selected as the active one. This is the candidate that will be used for actual exchange of media if and when its validated, until a higher priority candidate is validated. The active candidate will also be used to receive media from ICE-unaware peers. As such, it is RECOMMENDED that one be chosen based on the likelihood of that candidate to work with the peer that is being contacted. Unfortunately, it is difficult to ascertain which candidate that might be. As an example, consider a user within an enterprise. To reach non-ICE capable agents within the enterprise, a local candidate has to be used, since the enterprise policies may prevent communication between elements using a relay on the public network. However, when communicating to peers outside of the enterprise, a relayed candidate from a publically accessible STUN server is needed.

Indeed, the difficulty in picking just one address that will work is the whole problem that motivated the development of this specification in the first place. As such, it is RECOMMENDED that the active candidate be a relayed candidate from a STUN server providing public IP addresses in response to an Allocate request. Furthermore, ICE is only truly effective when it is supported on both sides of the session. It is therefore most prudent to deploy it to close-knit communities as a whole, rather than piecemeal. In the

example above, this would mean that ICE would ideally be deployed completely within the enterprise, rather than just to parts of it.

An additional consideration for selection of the active candidate is the switching of media stream destinations between the initial offer and the subsequent offer. If the active candidate pair in the initial offer is being validated, media will flow to that pair once it is validated. When the ICE checks complete and yield a higher priority candidate pair, media will begin to flow to it (there will also be an updated offer/answer exchange that changes the active candidate). This will result in a change in the destination of the media packets. This may also cause a different path for the media packets. That path might have different delay and jitter characteristics. As a consequence, the jitter buffers may see a glitch, causing possible media artifacts. If these issues are a concern, the initial offer MAY omit an active candidate. In such a case, an updated offer will need to be sent immediately when communicating with an ICE-unaware agent, setting an active candidate.

There may be transport-specific reasons for selection of an active candidate. In such a case, specifications defining usage of ICE with other transport protocols SHOULD document such considerations.

7.3 Encoding Candidates into SDP

For each candidate for a media stream, the agent includes a series of a=candidate attributes as media-level attributes, one for each component in the candidate. Each candidate has a unique identifier, called the candidate-id. The candidate-id MUST be chosen randomly and contain at least 24 bits of randomness (this does not mean that the candidate-id is 24 bits long; just that it has at least 24 bits of randomness). It is chosen only when the candidate is placed into the SDP for the first time; subsequent offers or answers within the same session containing that same candidate MUST use the same candidate-id used previously. 24 bits is sufficient because the candidate-id is not providing security (the much more random password is). It is needed only to prevent a possible simultaneous selection by two agents within a private network for the useful lifetime of the software or hardware.

Each component of the candidate has an identifier, called the component-id. The component-id is a sequence number. For each candidate, it starts at one, and increments by one for each component. As discussed below, ICE will perform connectivity checks such that, between a pair of candidates, checks only occur between transport addresses with the same component-id. As a consequence, if one candidate has three components, and it is paired with a candidate that has two, there will only be two transport address pairs and two

connectivity checks.

ICE will work without a standardized mapping between the components of a media stream and the numerical value of the component-id. This allows ICE to be used with media streams with multiple components without development of standards around such a mapping. However, a specific mapping has been defined in this specification for RTP - component-id 1 corresponds to RTP, and component-id of 2 corresponds to RTCP. Like the candidate-id, the component-id is assigned at the time the candidate is first placed into the SDP; subsequent offers or answers within the same session containing that same candidate MUST use the same component-id used previously.

The transport, addr and port of the a=candidate attribute (all defined in Section 12) are set to the transport protocol, unicast address and port of the tranport address. A Fully Qualified Domain Name (FQDN) for a host MAY be used in place of a unicast address. In that case, when receiving an offer or answer containing an FQDN in an a=candidate attribute, the FQDN is looked up in the DNS using an A or AAAA record, and the resulting IP address is used for the remainder of ICE processing. The qvalue is set to the priority of the candidate, and MUST be the same for all components of the candidate.

All of the candidates for a media stream share a password that is used for securing the STUN connectivity checks. Furthermore, the password for candidates for different media streams MAY be the same, or MAY be different. This password MUST be chosen randomly with 128 bits of randomness (though it can be longer than 128 bits). This password is contained in the a=ice-pwd attribute, present as a session or media level attribute. New passwords MUST be selected for each new session, even if the transport address from a previous session was being recycled.

The combination of candidate-id and component-id uniquely identify each transport address. As a consequence, each transport address has a unique identifier, called the tid. The tid is formed by concatenating the candidate-id with the component-id, separated by the colon (":"). The tid is not explicitly encoded in the SDP; it is derived from the candidate-id and component-id, which are present in the SDP. The usage of the colon as a separator allows the candidate-id and component-id to be extracted from the tid, since the colon is not a valid character for the candidate-id.

The tid gets combined, through further concatenation, with the tid of a transport address from the remote candidate (separated again by another colon) to form the username that is placed in the STUN checks between the peers. This allows the STUN message to uniquely identify the pairing whose connectivity it is checking. The tid is needed as

a unique identifier because the IP address within the candidate fails to provide that uniqueness as a consequence of NAT.

Consider agents A, B, and C. A and B are within private enterprise 1, which is using 10.0.0.0/8. C is within private enterprise 2, which is also using 10.0.0.0/8. As it turns out, B and C both have IP address 10.0.1.1. A sends an offer to C. C, in its answer, provides A with its transport addresses. In this case, thats 10.0.1.1:8866 and 8877. As it turns out, B is in a session at that same time, and is also using 10.0.1.1:8866 and 8877. This means that B is prepared to accept STUN messages on those ports, just as C is. A will send a STUN request to 10.0.1.1:8866 and 8877. However, these do not go to C as expected. Instead, they go to B. If B just replied to them, A would believe it has connectivity to C, when in fact it has connectivity to a completely different user, B. To fix this, tid takes on the role of a unique identifier. C provides A with an identifier for its transport address, and A provides one to C. A concatenates these two identifiers (with a colon between) and uses the result as the username in its STUN query to 10.0.1.1:8866. This STUN query arrives at B. However, the username is unknown to B, and so the request is rejected. A treats the rejected STUN request as if there were no connectivity to C (which is actually true). Therefore, the error is avoided.

An unfortunate consequence of the non-uniqueness of IP addresses is that, in the above example, B might not even be an ICE agent. It could be any host, and the port to which the STUN packet is directed could be any ephemeral port on that host. If there is an application listening on this socket for packets, and it is not prepared to handle malformed packets for whatever protocol is in use, the operation of that application could be affected. Fortunately, since the ports exchanged in SDP are ephemeral and ususally drawn from the dynamic or registered range, the odds are good that the port is not used to run a server on host B, but rather is the agent side of some protocol. This decreases the probability of hitting a port in-use, due to the transient nature of port usage in this range. However, the possibility of a problem does exist, and network deployers should be prepared for it. Note that this is not a problem specific to ICE; stray packets can arrive at a port at any time for any type of protocol, especially ones on the public Internet. As such, this requirement is just restating a general design guideline for Internet applications - be prepared for unknown packets on any port.

The active candidate, if there is one, is placed into the m/c lines of the SDP. For RTP streams, this is done by placing the RTP address and port into the c and m lines in the SDP respectively. If the agent is utilizing RTCP, it MUST encode its address and port using the a=rtcp attribute as defined in RFC 3605 [1]. If RTCP is not in

use, the agent MUST signal that using b=RS:0 and b=RR:0 as defined in RFC 3556 [6].

If there is no active candidate, the agent MUST include an a=inactive attribute. The RTP address and port in the m/c-line is inconsequential, since it won't be used.

Encoding of candidates may involve transport protocol specific considerations. There are none for UDP. However, extensions that define usage of ICE with other transport protocols SHOULD specify any special encoding considerations.

Once an offer or answer are sent, an agent MUST be prepared to receive both STUN and media packets on each candidate. As discussed in Section 7.13, media packets can be sent to a candidate prior to its promotion to active.

7.4 Forming Candidate Pairs

Once the offer/answer exchange has completed, both agents will have a set of candidates for each media stream. Each agent forms a set of candidate pairs for each media stream by combining each of its candidates with each of the candidates of its peer. Candidates can be paired up only if their transport protocols are identical. If an offer/answer exchange took place for a session comprised of an audio and a video stream, and each agent had two candidates per media stream, there would be 8 candidate pairs, 4 for audio and 4 for video. One agent can offer two candidates for a media stream, and the answer can contain three candidates for the same media stream. In that case, there would be six candidate pairs.

Each candidate has a number of components, each of which has a transport address. Within a candidate pair, the components themselves are paired up such that transport addresses with the same component ID are combined to form a transport address pair. Returning to the previous example, for each of the 8 candidate pairs, there would be two transport address pairs - one for RTP, and one for RTCP. If one candidate has more components than the other, those extra components will not be part of a transport address pair, won't be validated, and will effectively be treated as if they weren't included in the candidate pair in the first place.

The relationship between a candidate, candidate pair, transport address, transport address pair and component are shown in Figure 3. This figure shows the relationships as seen by the agent that owns the candidate with candidate ID "L". This candidate has two components with transport addresses A and B respectively. This candidate is called the native candidate, since it is the one owned

by the agent in question. The candidate owned by its peer is called the remote candidate. As the figure shows, there is a single candidate pair, and two components in each candidate. The native candidate has a candidate-id of "L", and the remote candidate has a candidate-id of "R". Since the two component-ids are 1 and 2, candidate "L" has two transport addresses with transport address IDs of "L:1" and "L:2" respectively. Similarly, candidate "R" has two transport addresses with transport address IDs of "R:1" and "R:2" respectively.

Furthermore, each transport address pair is associated with an ID, the transport address pair ID. This ID is equal to the concatenation of the tid of the native transport address with the tid of the remote transport address, separated by a colon. This means that the identifiers are seen differenly for each agent. For the agent that owns candidate "L", there are two transport address pairs. One contains transport address "L:1" and "R:1", with a transport address pair ID of "L:1:R:1". The other contains transport address "L:2" and "R:2", with a transport address pair ID of "L:2:R:2". For the agent that owns candidate "R", the identifiers for these two transport address pairs are reversed; it would be "R:1:L:1" for the first one and "R:2:L:2" for the second.

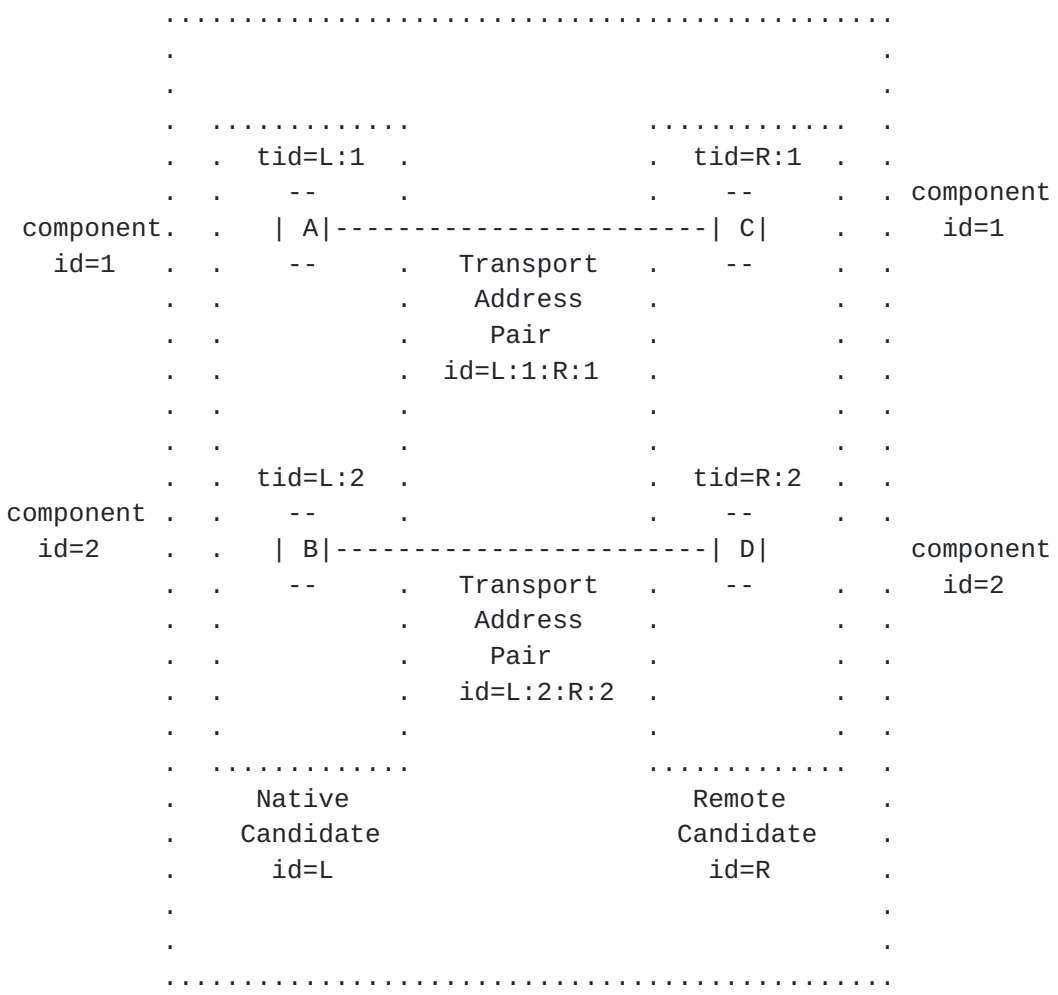

Candidate Pair

Figure 3

If a candidate pair was created as a consequence of an offer generated by an agent, then that agent is said to be the offerer of that candidate pair and all of its transport address pairs. Similarly, the other agent is said to be the answerer of that candidate pair and all of its transport address pairs. As a consequence, each agent has a particular role, either offerer or answerer, for each transport address pair. This role is important; when a candidate pair is to be promoted to active, the offerer is the one which performs the updated offer.

7.5 Ordering the Candidate Pairs

For the same reason that the STUN transactions during address gathering are paced at a rate of Ta transactions per second, so too

are the connectivity checks paced, also at a rate of Ta transactions per second. However, in order to rapidly converge on a valid candidate pair that is mutually desirable, the candidate pairs are ordered, and the checks start with the candidate pair at the top of the list. Rapid convergence of ICE depends on both the offerer and answerer coming to the same conclusion on the ordering of candidate pairs.

Recall that when each candidate is encoded into SDP, it contains a qvalue between 1 and 0, with 1 being the highest priority. Peer reflexive candidates, learned through the procedures described in Section 7.10 also have a priority between 0 and 1. For each media stream, the native candidates are ordered based on their qvalues, with higher q-values coming first. Amongst candidates with the same qvalue, they are ordered based on candidate ID, using reverse lexicographic order, where C1 is placed before C2, if C2 precedes C1 lexicographically. Lexicographic order can be viewed as a numerical ordering where each "digit" is actually a number in numerical base 256, with the mapping of characters to numerical value being defined by their ASCII encoding. For example, the candidate with candidate ID agD is greater than the candidate with ID ad7, and both of those are greater than the candidate with ID zz. Consequently, if these three candidates had equal q-values, they would be ordered as agD, ad7, zz - reverse of their lexicographic order.

The usage of a reverse lexicographic order is important; as discussed in Section 13, it allows peer-derived candidates to be preferred over native ones.

The result of these ordering rules will be an ordered list of candidates. The first candidate in this list is given a sequence number of 1, the next is given a sequence number of 2, and so on. This same procedure is done for the remote candidates. The result is that each candidate pair has two sequence numbers, one for the native candidate, and one for the remote candidate.

First, all of the candidate pairs for whom the smaller of the two sequence numbers equals 1 are taken first. Then, all of those for whom the smaller of the two sequence numbers equals 2 are taken next, and so on. Amongst those pairs that share the same value for their smaller sequence number, they are ordered by the larger of their two sequence numbers (smallest first). Amongst those pairs that share the same value for their smaller sequence number and the same value for their larger sequence number, the larger of the two candidate IDs in each pair are selected, and the pairs are lexicographically ordered in reverse by that candidate ID, largest first.

As an example, consider two agents, A and B. One offers two

candidates for a media stream with candidate IDs of "g9" and "88", with q-values of 1.0 and 0.8 respectively. The other answers with three candidates with candidate IDs of "h8", "65" and "kl", with q-values of 0.3, 0.2 and 0.1 respectively. The following table shows the rank ordering of the six candidate pairs. The column labeled "Max SN" is the larger of the two sequence numbers in the candidate pair, and "Min SN" is the minimum. The column labeled "Max Cand. ID" is the value of the larger of the two candidate IDs in the candidate pair.

Order	A Cand. ID	A Cand. q-value	A Cand. SN	B Cand. ID	B Cand. q-value	B Cand. SN	Max SN	Min SN	Max Cand. ID
1	g9	1.0	1	h8	0.3	1	1	1	h8
2	88	0.8	2	h8	0.3	1	2	1	h8
3	g9	1.0	1	65	0.2	2	2	1	g9
4	g9	1.0	1	k1	0.1	3	3	1	k1
5	88	0.8	2	65	0.2	2	2	2	88
6	88	0.8	2	k1	0.1	3	3	2	k1

This ordering is then modified slightly by taking the candidate pair corresponding to the active candidate, if there is one, and promoting it to the top of the list. To find this candidate pair, the agent looks for candidate pairs whose native and remote transport addresses match the native and remote transport addresses in the m/c-line. It is possible that multiple candidates match; this happens in the case where an agent obtained the same derived transport address from different local transport addresses. In such a case, the agent should pick one of the matching candidates.

Putting the active candidate at the top of the list allows it to be tested first. As discussed below, media is not sent until the corresponding candidate is verified, necessitating rapid verification of the active candidate. This modified ordering is called the candidate pair check ordering, since it reflects the order in which connectivity checks will be done. If there was no active candidate, the candidate pair check ordering and the candidate pair priority ordering will be identical.

Within each candidate pair there will be a set of transport address pairs, one for each component ID. Those pairs are ordered by component ID. The result is an absolute ordering of all transport address pairs for a media stream, sorted first by the order of their candidate pairs (with the exception of the active candidate), followed by the order of their component IDs. This ordering is

called the transport address pair check ordering.

Ordering of candidates may involve transport protocol specific considerations. There are none for UDP. However, extensions that define usage of ICE with other transport protocols SHOULD specify any special ordering considerations.

7.6 Performing the Connectivity Checks

Connectivity checks are a STUN usage defined in [13]. They are performed by sending peer-to-peer STUN Binding Requests. These checks result in a candidate progressing through a state machine that captures the progress of connectivity checks. The specific state machine and the procedures for the connectivity checks are specific to the transport protocol. This specification defines rules for UDP. Extensions to ICE that describe other transport protocols SHOULD describe the state machine and the procedures for connectivity checks.

The set of states visited by the offerer and answerer are depicted graphically in Figure 5

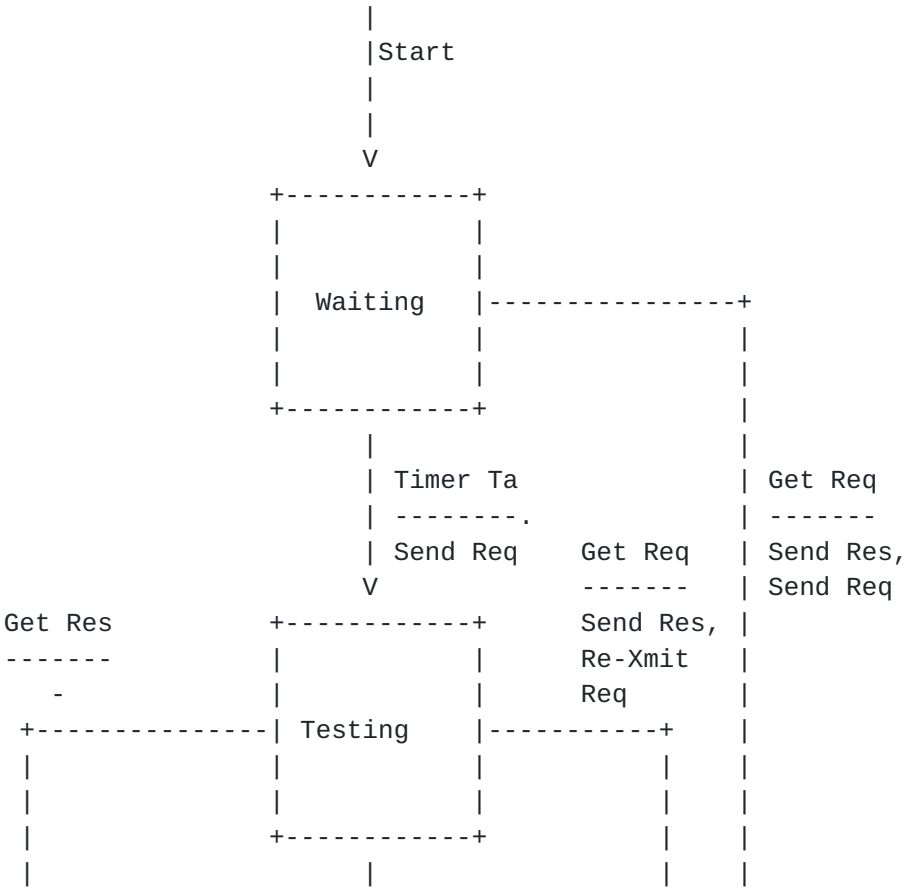

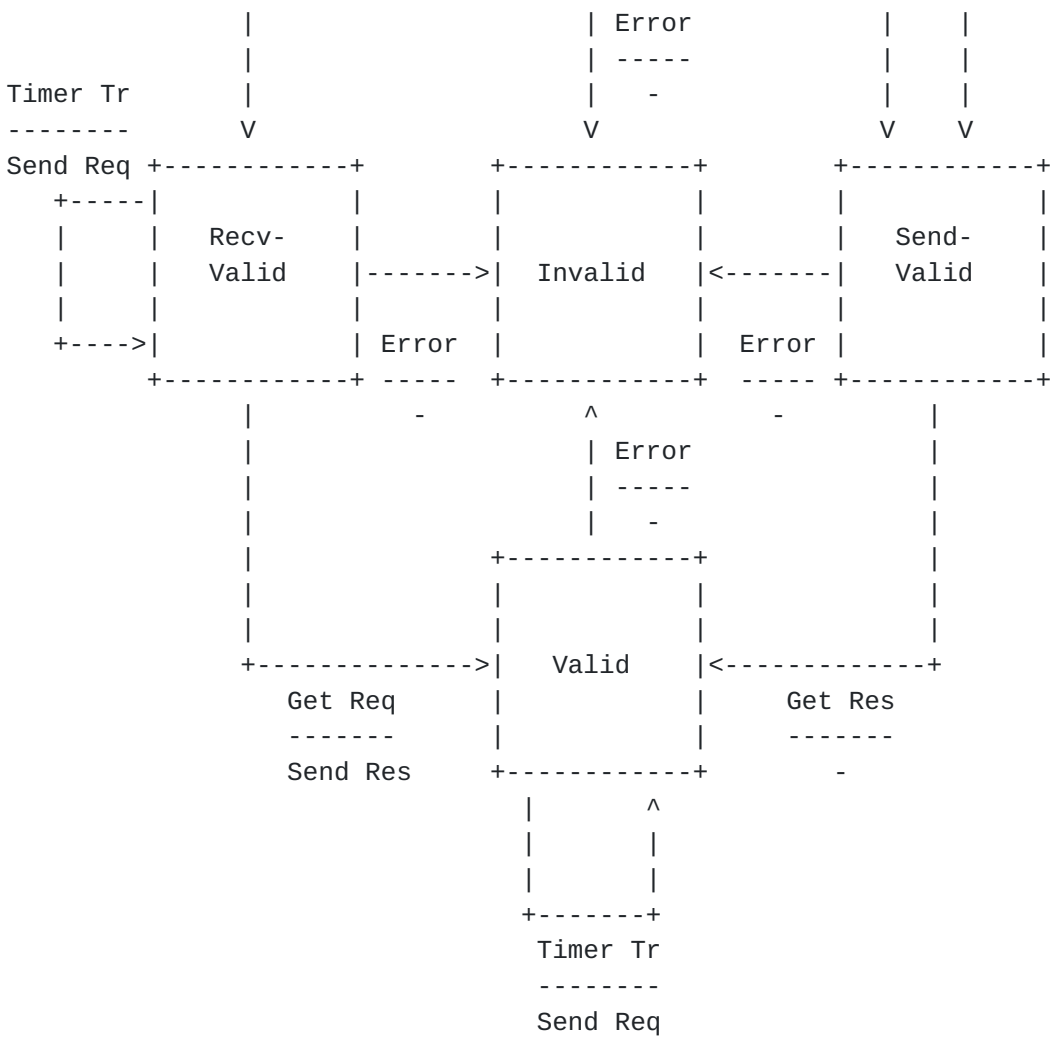

Figure 5

The state machine has six states - waiting, testing, Recv-Valid, Send-Valid, Valid and Invalid. Initially, all transport address pairs start in the waiting state. In this state, the agent waits for one of two events - a chance to send a Binding Request, or receipt of a Binding Request.

Since there is an instance of the state machine for each transport address pair, Binding Requests and responses need to be matched to the specific state machine for which they apply. This is done by computing the matching transport address pair for each Binding Request. This is done by examining the USERNAME of the incoming Binding Request. The USERNAME directly contains the transport address pair ID. Requests that are sent by an agent as part of the processing described here encode the transport address pair in the

USERNAME. Binding Responses are matched to their requests using the STUN transaction ID, and then mapped to the transport address pair from that.

Every Ta seconds, the agent starts a new connectivity check for a transport address pair. The check is started for the first transport address pair in the transport address pair check ordered list (which will be part of the active candidate) that is in the Waiting state. The state machine for this transport address pair is moved to the Testing state, and the agent sends a connectivity check using a STUN Binding Request, as outlined in Section 7.7. Once a STUN connectivity check begins, the processing of the check follows the rules for STUN. Specifically, retransmits of STUN requests are done as specified in [13], and furthermore, if a transaction fails and needs to be retried, that retry can happen rapidly, as described below. It doesn't "count" against the rate limit of 1/Ta checks per second. In addition, the keepalives that are generated for a valid pair do not count against the rate limit either. The rate limit applies strictly to the start of connectivity checks for a transport address pair that has been newly signaled through an offer/answer exchange.

In addition, if, while in the Waiting state, an agent receives a Binding Request matching that transport address pair, and this Binding Request generates a successful response, the transport address pair moves into the Send-Valid state, and the agent sends a connectivity check of its own using a STUN Binding Request, as outlined in Section 7.7. If the Binding Request didn't generate a success response, there is no change in state or generation of a Binding Request.

If, while in the Testing state, the agent receives a successful response to its STUN request, the transport address pair moves into the Recv-Valid state. In this state, the agent knows that packets can flow in both directions. However, its peer agent doesn't yet know that; all it knows is that it has been able to receive a packet. Thus, in this state, the agent awaits receipt of the Binding Request sent by its peer, as the response to that request is what informs its peer that packets can flow in both directions.

If, while in the Testing state, the agent receives a Binding Request matching that transport address pair, and this Binding Request generates a successful response, the transport address pair moves into the Send-Valid state. In addition, the agent retransmits a Binding Request for the transaction in progress. This helps speed up bidirectional connectivity verification when one agent is behind a symmetric NAT. If the Binding Request didn't generate a success response, there is no change in state or generation of a Binding

Request.

If, while in the Send-Valid state, the agent receives a successful response to its STUN request, the transport address pair moves to the Valid state. In this state, the agent knows that packets can flow in each direction. It also knows that its peer has sent it the STUN Request whose response will demonstrate to the peer that packets can flow in each direction.

If, while in the Recv-Valid state, the agent receives a STUN Binding Request from its peer that results in a successful response, the transport address pair moves into the Valid state. Receipt of a request whose response was not a successful one does not result in a change in state.

In any state, if the STUN transaction results in an error, the state machine moves into the invalid state. A STUN transaction produces an "error" based on the processing in Section 7.7, which indicates which STUN response codes constitute an error as far as ICE processing is concerned.

If a transport address pair is in the Recv-Valid or Valid state, an agent MUST generate a new STUN Binding Request transaction every Tr seconds. This transaction ensures that NAT bindings for the transport address pair remain open while the candidate is under consideration. The transaction is performed as outlined in Section 7.7. These transactions can also be used to keep the NAT bindings alive when the candidate is promoted to active, as described in Section 7.12. Tr SHOULD be configurable, and SHOULD default to 15 seconds. If the transaction results in an error, the state machine moves to the invalid state. This happens in cases where the NAT bindings expire (e.g., due to binding timeouts or NAT failures).

The candidate pair itself has a state, which is derived from the states of its transport address pairs. If at least one of the transport address pairs in a candidate pair is in the invalid state, the state of the candidate pair is considered to be invalid. If the candidate pair enters this state, an agent SHOULD move the state machines for all of the other transport address pairs in this candidate pair into the invalid state as well. This will ensure that connectivity checks never start for those transport address pairs. Furthermore, if checks are already in progress for one of those transport address pairs, the agent SHOULD cease them.

If all of the transport address pairs making up the candidate pair are Valid, the candidate pair is considered valid. If all of the transport address pairs making up the candidate pair are either Valid or Recv-Valid, and at least one is Recv-Valid, the candidate pair is

considered to be Recv-Valid. If all of the transport address pairs making up the candidate pair are either Valid or Send-Valid, and at least one is Send-Valid, the candidate pair is considered to be Send-Valid. If all of the transport address pairs in a candidate pair are in the Waiting state, the candidate pair is in the waiting state. If all of the transport address pairs in the candidate pair are either in the Waiting or Testing states, and at least one is in the Testing state, the state of the candidate pair is Testing. Otherwise, the state of the candidate pair is considered Indeterminate.

A candidate itself also has a state. If a candidate is present in at least one valid candidate pair, that candidate is said to be valid. If all of the candidate pairs containing that candidate are invalid, the candidate itself is invalid. Otherwise, the candidate's state is Indeterminate.

7.7 Sending a Binding Request for Connectivity Checks

An agent performs a connectivity check on a transport address pair by sending a STUN Binding Request from its native transport address, and sending it to the remote transport address. The meaning of "sending from its native transport address" depends on the type of transport protocol and the type of transport address (local, reflexive, or relayed). This specification defines the meaning for UDP. Specifications defining other transport protocols must define what this means for them.

For UDP-based local transport addresses, sending from the local transport address has the meaning one would expect - the request is sent such that the source IP address and port equal that of the local transport address. For reflexive ransport addresses, it is sent by sending from the associated local transport address used to derive that reflesive address. For relayed transport addresses, it is sent by using STUN mechanisms to send the request through the STUN relay (using the Send request). Sending the request through the STUN relay server neccesarily requires that the request be sent from the client, using the local transport address used to derive the relayed transport address.

The Binding Request sent by the agent MUST contain the USERNAME attribute. This attribute MUST be set to the transport address pair ID of the corresponding transport address pair as seen by its peer. Thus, for the first transport address pair in Figure 3, if the agent on the left sends the STUN Binding Request, the USERNAME will have the value R:1:L:1. If the agent on the right sends the STUN Binding Request, the USERNAME will have the value L:1:R:1. To be clear, the USERNAME that is used is NOT the one seen locally, but rather the one as seen by its peer. The request SHOULD contain the MESSAGE-

INTEGRITY attribute, computed according to [13]. The key used as input to the HMAC is the password provided by the peer for this remote transport address. This password will be identical for all remote transport addresses for the same media stream.

Note that all ICE implementations are required to be compliant to [13], as opposed to the older [16]. Consequently, all connectivity checks will contain the magic cookie in the STUN header, and cause the STUN server embedded in each ICE implementation to include XOR-MAPPED-ADDRESS attributes in the response, rather than MAPPED-ADDRESS.

The STUN transaction will generate either a timeout, or a response. If the response is a 420, 500, or 401, the agent should try again as described in [13] (as mentioned above, it need not wait Ta seconds to try again). Either initially, or after such a retry, the STUN transaction might produce a non-recoverable failure response (error codes 400, 430, 431, or 600) or a failure result inapplicable to this usage of STUN and thus unrecoverable (432, 433). If this happens, an error event is generated into the state machine, and the transport address pair enters the invalid state.

If the STUN transaction times out, the client SHOULD NOT retry. The only reason a retry might succeed is if there was severe packet loss during the duration of the check, or the answer was significantly delayed, also due to packet loss. However, STUN Binding Request transactions run for 9.5 seconds, which is well beyond the typical tolerance for a session establishment. The retries come with a penalty of additional traffic, which can be used to launch DoS attacks Section 13.4.2. The only reason to not follow the SHOULD NOT is if the agent has adjusted the STUN transaction timers to be more aggressive.

If the Binding Response is a 200, the agent SHOULD check for the MESSAGE-INTEGRITY attribute and verify it, as discussed in [13]. Indeed, this check SHOULD be done for all responses. This will result in the response being discarded (eventually leading to a timeout), if the integrity check fails.

7.8 Receiving a Binding Request for Connectivity Checks

As a result of providing a list of candidates in its offer or answer, an agent will receive STUN Binding Request messages. An agent MUST be prepared to receive STUN Binding Requests on each local transport address from the moment it sends an offer or answer that contains a candidate with that local transport address. Similarly, it MUST be prepared to receive STUN Binding Requests on a local transport address the moment it sends an offer or answer that contains a

reflexive or relayed candidate derived from a local candidate with that local transport address. It can cease listening for STUN messages on that local transport address after sending an updated offer or answer which does not include any candidates with transport addresses that are equal to or derived from that local transport address.

As discussed in [13], since the username and password for STUN requests are exchanged through another mechanism - here, ICE - the Shared Secret Request mechanism is not needed and need not be implemented by agents that provide the connectivity check usage.

One of the candidates may be in use as the active candidate, or may become promoted to the active candidate in the next offer/answer exchange as a consequence of a successful validation. In either case, both media and STUN packets will be sent to the transport addresses comprising that candidate, causing both to receive on their associated local transport addresses. The agent MUST be able to disambiguate them. This is done trivially by looking for the STUN magic cookie as the value of the second 32-bit word in the packet. If present, it identifies a STUN packet.

Processing of the Binding Request proceeds in two steps. The first is generation of the response, and the second ICE-specific processing. Generation of the response follows the general procedures of [13]. The USERNAME is considered valid if one of the candidate IDs sent in an offer or answer is a prefix of the USERNAME (this will always be the case, even for peer reflexive candidates). The password associated with that candidate ID is used to verify the MESSAGE-INTEGRITY attribute, if one was present in the request. If the USERNAME was not valid, the agent generates a 430. Otherwise, the success response will include the XOR-MAPPED-ADDRESS attribute, which is used for learning new candidates, as described in Section 7.10. The XOR-MAPPED-ADDRESS attribute is constructed using the source IP address and port of the Binding Request. For Binding Requests received over relayed transport addresses, this MUST be the source IP address and port of the Binding Request when it arrived at the relay, prior to forwarding towards the agent. That source transport address will be present in the REMOTE-ADDRESS attribute of a STUN Data Indication message, if the Binding Request was delivered through a Data Indication. If the Binding Request was not encapsulated in a Data Indication, that source address is equal to the current active destination for the STUN relay session.

The ICE processing involves changes to the state machine for a transport address pair. This processing cannot be done until the initial offer/answer exchange has completed. As a consequence, if the oferrer received a Binding Request that generated a success

response, but had not yet received the answer to its offer, it waits for the answer, and when it arrives, then performs the ICE processing.

The agent takes the entire contents of the USERNAME, and compares them against the transport address pair identifiers as seen by that agent for each transport address pair. If there is no match, nothing is done - this should never happen for compliant implementations. If there is a match, the resulting transport address pair is called the matching transport address pair. The state machine for the matching transport address pair is then updated based on the receipt of a STUN Binding Request, and the resulting actions described in Section 7.6 are undertaken.

An agent will continue to receive periodic STUN connectivity checks on a local transport address as long as it had listed that transport address, or one derived from it, in an a=candidate attribute in its most recent offer or answer, the state machine for that transport address is in the Recv-Valid or Valid states, and the transport address is for UDP. Whether STUN keepalives are used for other transport protocols is defined by the specifications for that transport protocol. The agent processes any such transactions according to this section. It is possible that a transport address pair that was previously valid may become invalidated as a result of a subsequent failed STUN transaction.

7.9 Promoting a Candidate to Active

As a consequence of the connectivity checks, each agent will change the states for each transport address pair, and consequently, for the candidate pairs. When a candidate pair becomes valid, and the agent is in the role of offerer for that candidate pair, the agent follows the logic in this section. The rules only apply to the offerer of a candidate pair in order to eliminate the possibility of both agents simultaneously offering an update to promote a candidate to active.

If this candidate pair is the first one in the candidate pair priority ordered list, the agent SHOULD send an updated offer as described in Section 7.11.1. If this candidate pair is not the first on that list, but it is the first on the candidate pair check ordered list, it means that this candidate pair is the active one, and its connectivity has been verified. This is good news; the currently active candidate is working. Media can now flow as described in Section 7.13 (media will never flow prior to validation). However, no updated offer is sent at this time.

If this candidate pair is not the first on the candidate pair priority ordered list or the candidate pair check ordered list, and

the wait-state timer has not yet been set, the agent sets this timer to Tws seconds. Tws SHOULD be configurable, and SHOULD have a default of 100ms. This timer allows for a higher priority connectivity check to complete, in the event its STUN Binding Request was lost or delayed in the network. If, prior to the wait-state timer firing, another connectivity check completes and a candidate pair is validated, there is no need to reset or cancel the timer. Once the timer fires, the agent SHOULD issue an updated offer as described in Section 7.11.1.

In addition, in order to speed up ICE processing, once the agent has determined the candidate that is to be promoted, it will send and receive media using that candidate in expectation of an updated offer. This is discussed in Section 7.13.

7.10 Learning New Candidates from Connectivity Checks

ICE makes use of reflexive addresses, which are addresses that inform an agent of its transport address as seen by another host. An initial offer or answer generated by an agent includes server reflexive addresses, which are learned from a configured or discovered STUN server in the network. However, the connectivity checks themselves can inform an agent of reflexive addresses, and in particular, ones that are reflexive towards its peer. These are called peer reflexive candidates. A new peer reflexive candidate is typically observed when two agents are separated by a NAT with the address-dependent or address and port dependent mapping properties [37]. When the agent behind such a NAT sends a Binding Request to the other agent (assuming it is reachable), the NAT will create a new mapping for this Binding Request. Because STUN and the media packets are sent on the same port, regardless of the filtering properties of the NAT (whether endpoint independent, address dependent, or address and port dependent), this reflexive address can be used by the peer for sending STUN and media packets back towards the agent.

To obtain and use these peer reflexive transport addresses, ICE agents perform additional processing on the receipt of STUN Binding Requests and responses, beyond the logic described in Section 7.7 and Section 7.8. This logic is described below.

7.10.1 On Receipt of a Binding Request

When a STUN Binding Request is received which generates a success response, that Binding Request would have been associated with a matching transport address pair and corresponding candidate pair. The source IP and port of this Binding Request are compared to the IP address and port of the remote transport address in the matching transport address pair. Note that, in this case, we are comparing

actual IP addresses and ports - not tids. In addition, if the Binding Request arrived through a relayed transport address, the source IP and port of this binding request used for the comparison are those in the Binding Request when it arrived at the relay, prior to forwarding towards the agent. That source transport address will be present in the REMOTE-ADDRESS attribute of a STUN Data Indication message, if the Binding Request were delivered through a Data Indication. If the Binding Request was not encapsulated in a Data Indication, that source address is equal to the current active destination for the STUN relay session.

The comparison of the source IP and port of the Binding Request and the IP address and port of the remote transport address in the matching transport address pair may indicate inequality. In that case, the source IP and port of the Binding Request (and again, for relayed transport address, this refers to the source IP address and port of the packet when it arrived at the relay) are compared to the IP address and ports across the transport address pairs in *all* remote candidates. If there is a match to another remote candidate (called the alternate remote candidate), this is not a new candidate; however, the Binding Request has effectively helped validate the alternate remote candidate. The agent SHOULD select the candidate pair corresponding to the combination of the alternate remote candidate and the native candidate from the original matching candidate pair. A "Get Req" event is passed to the state machine for that candidate pair. Consequently, if this candidate pair was in the Waiting state, a connectivity check will be generated for it.

If, when the source IP and port of the STUN packet, when compared against all remote candidates, was not a match to any of them, it means that the source IP and port might represent another valid remote transport address - a peer derived one.

To use it, that address needs to be associated with a candidate (called a peer-derived candidate). In this case, however, the candidate isn't signaled through an offer/answer exchange; it is constructed dynamically from information in the STUN request. Like all other candidates, the peer-derived candidate has a candidate ID. The candidate ID is derived from the candidate IDs of the matching candidate pair. In particular, the candidate ID is constructed by concatenating the remote candidate ID with the native candidate ID (without the colon). The password for the new candidate equals that of the remote candidate ID in the matching candidate pair (note that, this password would be the same for all remote candidates for the same media line).

On receipt of a STUN Binding Request whose source IP and port don't match the transport address in any remote candidate, the agent

constructs the candidate ID that represents the peer reflexive candidate, and checks to see if that candidate exists. It may already exist if it had been constructed as a consequence of a previous application of this logic on receipt of a Binding Request for a different transport address pair of the same candidate pair. If there is not yet a peer reflexive candidate with that candidate ID, the agent creates it, and assigns it the newly computed candidate ID. The priority of the peer-derived candidate MUST be set to the priority of its generating candidate - the remote candidate in the matching transport address pair. Note that, at this time, the peer derived candidate has no transport addresses in it.

Newly created or not, the agent extracts the component ID from the matching transport address pair, and sees if a transport address with that same component ID exists in the peer reflexive candidate. If not (and it shouldn't), the agent adds a transport address to the peer reflexive candidate. This transport address is equal to the source IP address and port from the incoming STUN Binding Request (and in the case of a relayed transport address, the one seen by the relay). It is assigned the component ID equal to the component ID in the matching transport address pair. This transport address will have a tid, equal to the concatenation of the candidate ID for this new candidate, and the component ID, separated by a colon.

The peer reflexive candidate becomes usable once the number of transport addresses in it equals the transport address pair count of the candidate pair from which it is derived. Initially, the peer reflexive candidate will start with a single transport address. More are added as the connectivity checks for the original candidate pair take place. Once the peer reflexive candidate becomes usable, it has to be paired up with native candidates. However, unlike the procedures of Section 7.5, which pair up each remote candidate with each native candidate, this peer reflexive candidate is only paired up with the native candidate from the candidate pair from which it was derived. This creates a new candidate pair, and a set of new transport address pairs.

Recall that, for each candidate pair, one agent plays the role of offerer, and the other of answerer. For a peer-reflexive candidate, the role is identical to that of its generating candidate.

Figure 6 provides a pictorial representation of the peer reflexive candidate (the one with id=RL) and its pairing with the native candidate with id L. The candidate with ID R is referred to as the generating candidate. The peer reflexive candidate is effectively an alternate for that generating candidate, but is only paired with a specific native candidate. Note that, for a particular generating candidate, there can be many peer derived candidates, up to one for

each native candidate.

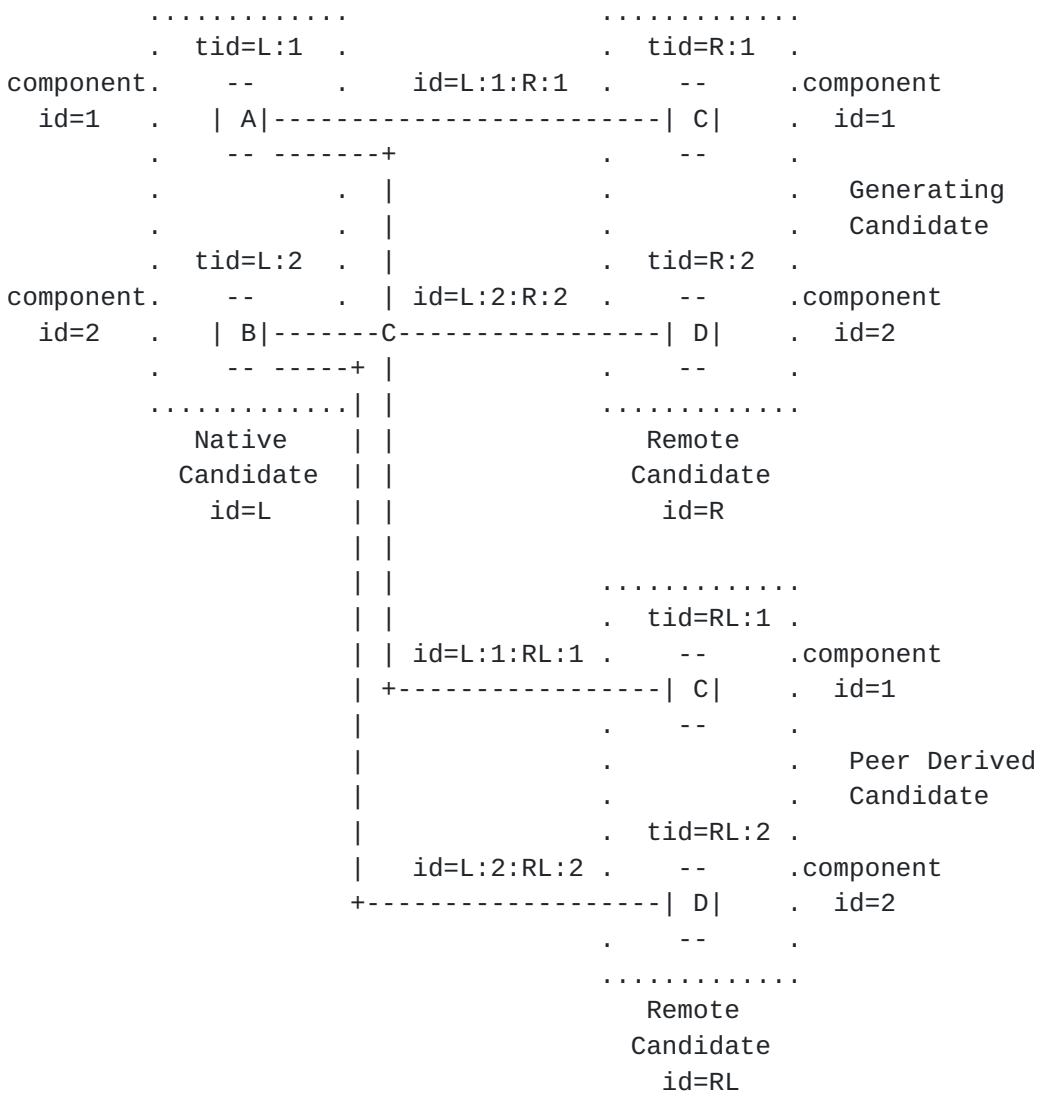

Figure 6

The new transport address pairs have a state machine associated with them. The state that is entered, and actions to take as a consequence, are specific to the transport protocol. For UDP, the procedures are defined here. Extensions that define processing for other transport protocols SHOULD describe the behavior.

For UDP, the state machine enters the Send-Valid state. Effectively, the Binding Request just received "counts" as a validation in this direction, even though it was formally done for a different candidate pair. In addition, the agent SHOULD generate a Binding Request for each transport address in this new candidate pair, as described in

Section 7.7. The transport address pairs are inserted into the ordered list of pairs based on the ordering described in Section 7.5 and processing follows the logic described in Section 7.6.

7.10.2 On Receipt of a Binding Response

The procedures on receipt of a Binding Response are nearly identical to those for receipt of a Binding Request as described above.

When a successful STUN Binding Response is received, it will be associated with a matching transport address pair and corresponding candidate pair. This matching is done based on comparison of candidate IDs. The reflexive transport address from the Binding Response is compared to the IP address and port of the native transport address in the matching transport address pair. Note that, in this case, we are comparing actual IP addresses and ports - not tids. These may not match if there was a NAT between the two agents. If they do not match, the reflexive transport address is compared to the IP address and ports across the transport address pairs in *all* native candidates. If there is a match to another native candidate (called the alternate native candidate), this is not a new candidate; however, the Binding Response has effectively helped validate the alternate native candidate. The agent SHOULD select the candidate pair corresponding to the combination of the alternate native candidate and the remote candidate from the original matching candidate pair. If the candidate pair is in the Waiting state, it moves directly to the Recv Valid state.

If, when the reflexive transport address, when compared against all native candidates, was not a match to any of them, it means that the reflexive transport address might represent another valid native transport address - a peer derived one.

To use it, that address needs to be associated with a candidate. In this case, however, the candidate isn't signaled through an offer/answer exchange; it is constructed dynamically from information in the STUN response. Such a candidate is called a peer reflexive candidate. Like all other candidates, the peer reflexive candidate has a candidate ID. The candidate ID is derived from the candidate IDs of the matching candidate pair. In particular, the candidate ID is constructed by concatenating the native candidate ID with the remote candidate ID (without the colon). The password for the new candidate equals that of the native candidate ID in the matching candidate pair.

On receipt of a STUN Binding Response whose reflexive transport address didn't match the transport address in any native candidate, the agent constructs the candidate ID that represents the peer

reflexive candidate, and checks to see if that candidate exists. It may already exist if it had been constructed as a consequence of a previous application of this logic on receipt of a Binding Response for a different transport address pair of the same candidate pair. If there is not yet a peer derived candidate with that candidate ID, the agent creates it, and assigns it the newly computed candidate ID. The priority of the new candidate MUST be set to the priority of the generating candidate - the native candidate in the matching transport address pair. Note that, at this time, the peer derived candidate has no transport addresses in it.

Newly created or not, the agent extracts the component ID from the matching transport address pair, and sees if a transport address with that same component ID exists in the peer reflexive candidate. If not (and it shouldn't), the agent adds a transport address to the peer reflexive candidate. This transport address is equal to the reflexive transport address from the STUN Binding Response. It is assigned the component ID equal to the component ID in the matching transport address pair. This transport address will have a tid, equal to the concatenation of the candidate ID for this new candidate, and the component ID, separated by a colon.

The peer-derived candidate becomes usable once the number of transport addresses in it equals the transport address pair count of candidate pair from which it is derived. Initially, the peer-derived candidate will start with a single transport address. More are added as the connectivity checks for the original candidate pair take place. Once the peer-derived candidate becomes usable, it has to be paired up with remote candidates. However, unlike the procedures of Section 7.5, which pair up each remote candidate with each native candidate, the peer-derived candidate is only paired up with the remote candidate from the matching candidate pair. This creates a new candidate pair, and a set of new transport address pairs.

Recall that, for each candidate pair, one agent plays the role of offerer, and the other of answerer. For a peer-reflexive candidate, the role is identical to that of its generating candidate.

The new transport address pairs have a state machine associated with them. The state that is entered, and actions to take as a consequence, are specific to the transport protocol. For UDP, the procedures are defined here. Extensions that define processing for other transport protocols SHOULD describe the behavior.

For UDP, the state machine enters the Recv-Valid state. Effectively, the Binding Response just received "counts" as a validation in this direction, even though it was formally done for a different candidate pair. The transport address pairs are inserted into the ordered list

of pairs based on the ordering described in Section 7.5, and processing follows the logic described in Section 7.6.

7.11 Subsequent Offer/Answer Exchanges

An agent MAY issue an updated offer at any time. This updated offer may be sent for reasons having nothing to do with ICE processing (for example, the addition of a video stream in a multimedia session), or it may be due to a change in ICE-related parameters. For example, if an agent acquires a new candidate after the initial offer/answer exchange, it may seek to add it.

However, agents SHOULD follow the logic described in Section 7.9 to determine when to send an updated offer as a consequence of promoting a candidate to active.

If there are any aspects of this processing that are specific to the transport protocol, those SHOULD be called out in ICE extensions that define operation with other transport protocols. There are no additional considerations for UDP.

7.11.1 Sending of a Subsequent Offer

The offer MAY contain a new active candidate in the m/c line. This candidate SHOULD be the native candidate from the highest candidate pair in the candidate pair priority ordered list whose state is Valid. If there are no candidate pairs in this state, the highest one whose state is Send-Valid or Recv-Valid SHOULD be used. If there are no candidate pairs in these states, the candidate pair that is most likely to work with this peer, as described in Section 7.2, SHOULD be used. The candidate is encoded into the m/c line in an updated offer as described in Section 7.3. Note that, while peer-derived candidates never appear in a=candidate attributes (only their generating candidates appear there), a peer-derived candidate can appear in the m/c line if it has been selected for usage for media.

If the candidate pair whose native candidate was encoded into the m/c-line was Valid, Send-Valid or Recv-Valid, the agent MUST include an a=remote-candidate attribute into the offer. This attribute MUST contain the candidate ID of the remote candidate in the candidate pair. It is used by the recipient of the offer in selecting its candidate for the answer.

The meaning of a=candidate attributes within a subsequent offer have the same meaning as they do in an initial offer. They are a request for the peer to attempt (or continue to attempt if the candidate was provided previously) a connectivity check using STUN from each of its own candidates. When an updated offer is sent, there are several

dispositions regarding the candidates:

retained: A candidate is retained if the candidate ID for the candidate is included in the new offer, and matches the candidate ID for a candidate in the previous offer or answer from the agent. In this case, all of the information about the candidate - its qvalue and components, and the IP addresses, ports, and transport protocols of its components, MUST be the same as the previous offer or answer from the agent. If the agent wants to change them, this is accomplished by changing the candidate ID as well. That will have the effect of removing the old candidate and adding a new one with the updated information.

removed: A candidate is removed if its candidate ID appeared in a previous offer or answer, and that candidate ID is not present in the new offer.

added: A candidate is added if its candidate ID appeared in the new offer, but was not present in a previous offer or answer from that agent.

The following rules are used to determine the disposition of the each of the current native candidates in the new offer:

- If a candidate is invalid, and all peer reflexive candidates generated from it are invalid as well, it SHOULD be removed.

- If the candidate in the m/c-line is valid, all other candidates SHOULD be removed. This has the effect of stopping connectivity checks of other candidates. This SHOULD would not be followed if an agent wanted to keep a candidate ready for usage should, for some reason, the active candidate later become invalid.

- If the candidate in the m/c-line is valid, and it is not peer reflexive, that candidate MUST be retained. If the candidate in the m/c-line is peer reflexive, its generating candidate MUST be retained, even if it is itself invalid.

- If the candidate in the m/c-line has not been validated, all other candidates that are not invalid, or candidates for whom their derived candidates are not invalid, SHOULD be retained.

- Peer reflexive candidates MUST NOT be added; they continue to be used as long as their generating candidate was retained. Peer derived candidates are learned exclusively through the STUN connectivity checks.

A new candidate MAY be added. This can happen when the candidate is

a new one, learned since the previous offer/answer exchange, and it has a higher priority than the currently active candidate. It can also occur when an agent wishes to restart checks for a transport address it had tried previously. Effectively, changing the candidate ID value in an updated offer will "restart" connectivity checks for that candidate.

If a candidate is removed, the agent takes the following steps once the offer is sent:

1. The agent eliminates any candidate pairs whose native candidate equalled the candidate that was removed. Equality is based on comparison of candidate IDs.

2. The agent eliminates any candidate pairs that had a native candidate that is a peer reflexive candidate generated from the candidate that was removed.

3. The candidate pairs that are eliminated are removed from the candidate pair priority ordered list and candidate pair check ordered list. As a consequence of this, if connectivity checks had not yet begun for the candidate pair, they won't.

4. If connectivity checks were already in progress for transport addresses in a candidate pair that was removed, the agent SHOULD immediately terminate them. No further retransmissions take place, and no further transactions from that candidate will be made.

5. If the removed candidate was a relayed candidate, the agent SHOULD de-allocate its transport addresses from the STUN relay if it is not using those resources elswhere. If a local candidate was removed, and all of its derived candidates were also removed (including any peer reflexive candidates), local operating system resources for each of the transport addresses in the local candidate SHOULD be de-allocated, as long as it is not using those resources elsewhere. The resources may be in use elsewhere if they were included in an initial offer which generated multiple answers (as can happen with SIP forking). In such a case, a subsequent offer which removes the candidate will not imply its removal with the other branches; each becomes a separate offer/answer relationship.

Subsequent offers MUST contain a=ice-pwd attributes that specify the password for the candidates for each media stream. The password for the candidates for a particular media stream SHOULD have the same value as in previous offers. However, an agent MAY change it if, for some reason, the agent believes that the password may have been

compromised. Note that it is permissible to use a session-level attribute in one offer, and in a subseqeunt offer, provide the same password as a media-level attribute. This is not a change in the password; merely a change in its representation. An agent MUST be prepared to receive connectivity checks that use either the new or old password until Tpw seconds after it receives the answer. Tpw SHOULD be configurable, and SHOULD default to 2 seconds.

7.11.2 Receiving the Offer and Sending an Answer

To generate the answer, the answerer has to decide which transport addresses to include in the m/c line, and which to include in candidate attributes.

The first step in the process is to look for the a=remote-candidate attribute in the offer. The a=remote-candidate exists to eliminate a race condition between the updated offer and the response to the STUN Binding Request that moved a candidate into the Valid state. This race condition is shown in Figure 7. On receipt of message 5, agent A can move its transport address pair state machine into the Valid state. It sends a STUN response to the request (message 6), but this is lost. Agent A proceeds with an updated offer (message 7), which is received at agent B. As far as agent B is concerned, the transport address pair is still in the Send-Valid state. It will move into the Valid state only on receipt of the STUN response in message 10. Thus, upon receipt of the offer, agent B cannot determine which candidate to include in its answer. To eliminate this condition, the identity of the validated candidate is included in the offer itself. Note, however, that the answerer will not send media until it has received this STUN response.

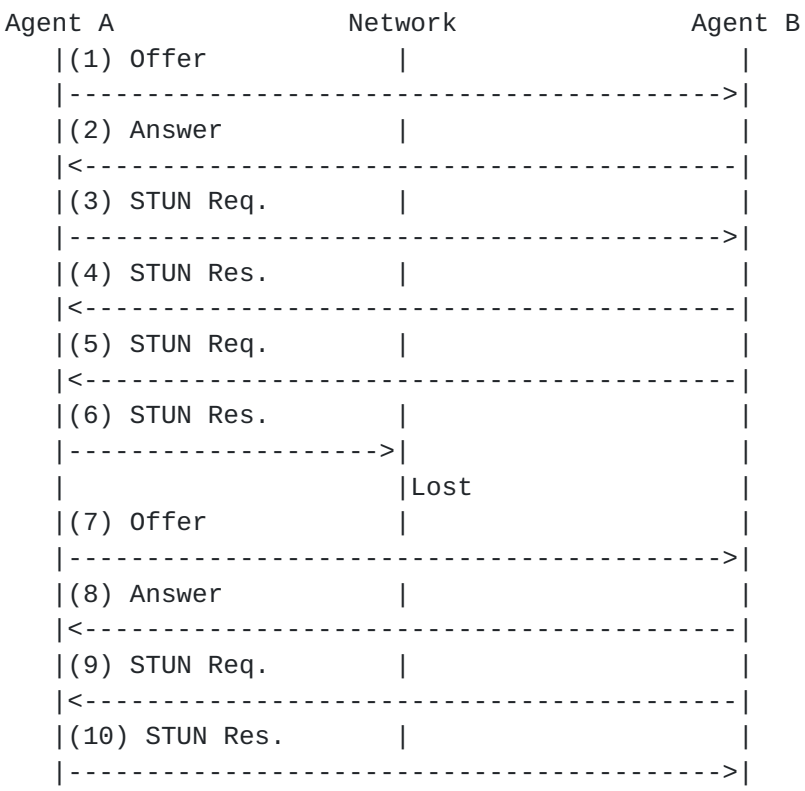

Figure 7

If the a=remote-candidate attribute is present, the agent examines the transport addresses in the m/c-line of the offer. It compares these with the transport addresses in the remote candidates of all candidate pairs. If there is at least one match, the agent compares the native candidate ID of each matching pair with the value of the a=remote-candidate attribute. If there is a match, that candidate pair is selected. For each transport address pair in that candidate pair, if the state of the transport address pair is Send-Valid, the agent considers the state to be Valid just for the purpose of selecting the m/c-line as discussed in the paragraph below. The actual state MUST remain Send-Valid. This is necessary to prevent against DoS attacks.

Rules for choosing transport addresses for the m/c-line are as follows. The agent examines the transport addresses in the m/c-line of the offer. It compares these with the transport addresses in the remote candidates of candidate pairs whose states are Valid. If there is a matching candidate pair in that state, the pair with the highest priority MUST be chosen, and the native candidate from that pair used as the active candidate. If there were no matching candidate pairs in the Valid state, the candidate that is most likely to work with this peer, as described in Section 7.2, SHOULD be used.

Like the offerer, the answerer can decide, for each of its candidates, whether they are retained or removed. The same rules defined in Section 7.11.1 for determining their disposition apply to the answerer. Similarly, if a candidate is removed, the same rules in Section 7.11.1 regarding removal of canididate pairs and freeing of resources apply.

Once the answer is sent, the answerer will have the set of native and remote candidates before this offer/answer exchange, and the set of native and remote candidates afterwards. A peer derived candidate continues to be used as long as its generating parent continues to be used. The agent then pairs up the native and remote candidates which were added or retained. This leads to a set of current candidate pairs.

If a candidate pair existed previously, but as a consequence of the offer/answer exchange, it no longer exists, the agent takes the following steps:

1. The candidate pair is removed from the candidate pair priority ordered list and candidate pair check ordered list. As a consequence of this, if connectivity checks had not yet begun for the candidate pair, they won't.

2. If connectivity checks were already in progress for that candidate pair, the agent SHOULD immediately terminate any STUN transactions in progress from that candidate. No further retransmissions take place, and no further transactions from that candidate will be made.

3. If the agent receives a STUN Binding Request for that candidate pair, the agent SHOULD generate a 430 response.

If a candidate pair existed previously, and continues to exist, no changes are made; any STUN transactions in progress for that candidate pair continue, and it remains on the candidate pair priority ordered list and candidate pair check ordered list.

If a candidate pair is new (because either its native candidate is new, or its remote candidate is new, or both), the agent takes the role of answerer for this candidate pair. The new candidate pair is inserted into the candidate pair priority ordered list and candidate pair check ordered list. STUN connectivity checks will start for them based on the logic described in Section 7.6.

7.11.3 Receiving the Answer

Once the answer is received, the answerer will have the set of native

and remote candidates before this offer/answer exchange, and the set of native and remote candidates afterwards. It then follows the same logic described in Section 7.11.2, pairing up the candidate pairs, removing ones that are no longer in use, and beginning of processing for ones that are new.

7.12 Binding Keepalives

Once a candidate is promoted to active, and media begins flowing, it is still necessary to keep the bindings alive at intermediate NATs for the duration of the session. Normally, the media stream packets themselves (e.g., RTP) meet this objective. However, several cases merit further discussion. Firstly, in some RTP usages, such as SIP, the media streams can be "put on hold". This is accomplished by using the SDP "sendonly" or "inactive" attributes, as defined in RFC 3264 [4]. RFC 3264 directs implementations to cease transmission of media in these cases. However, doing so may cause NAT bindings to timeout, and media won't be able to come off hold.

Secondly, some RTP payload formats, such as the payload format for text conversation [36], may send packets so infrequently that the interval exceeds the NAT binding timeouts.

Thirdly, if silence suppression is in use, long periods of silence may cause media transmission to cease sufficiently long for NAT bindings to time out.

To prevent these problems, ICE implementations MUST continue to list their active candidate in a=candidate lines for UDP-based media streams. As a consequence of this, STUN packets will be transmitted periodically independently of the transmission (or lack thereof) of media packets. This provides a media independent, RTP independent, and codec independent solution for keeping the NAT bindings alive.

If an ICE implementation is communciating with one that does not support ICE, keepalives MUST still be sent. Indeed, these keepalives are essential even if neither endpoint implements ICE. As such, this specification defines keepalive behavior generally, for endpoints that support ICE, and those that do not.

All endpoints MUST send keepalives for each media session. These keepalives MUST be sent regardless of whether the media stream is currently inactive, sendonly, recvonly or sendrecv. The keepalive SHOULD be sent using a format which is supported by its peer. ICE endpoints allow for STUN-based keepalives for UDP streams, and as such, STUN keepalives MUST be used when an agent is communicating with a peer that supports ICE. An agent can determine that its peer supports ICE by the presence of the a=candidate attributes for each

media session. If the peer does not support ICE, the choice of a packet format for keepalives is a matter of local implementation. A format which allows packets to easily be sent in the absence of actual media content is RECOMMENDED. Examples of formats which readily meet this goal are RTP No-Op [31] and RTP comfort noise [26].

STUN-based keepalives will be sent periodically every Tr seconds as a consequence of the rules in in Section 7.7. If STUN keepalives are not in use (because the peer does not support ICE), an agent SHOULD ensure that a media packet is sent every Tr seconds. If one is not sent as a consequence of normal media communications, a keepalive packet using one of the formats discussed above SHOULD be sent.

7.13 Sending Media

When an agent receives an offer and sends an answer, or when it receives an answer to an offer it sent, it begins connectivity checks. These checks will include validation of the active candidate pair, if there was one. An agent SHOULD NOT send media on the active candidate pair until that candidate pair has reached the Valid or Recv-Valid state. This is to help prevent a denial-of-service attack, described in Section 13. Once the active candidate pair reaches the Valid or Recv-Valid state, an agent MAY start sending media to that candidate pair.

However, offer/answer exchanges are used with protocols, like SIP, which require media to be sent "early", from the answerer to the offer, prior to completion of the initial offer/answer exchange. It is highly desirable (and sometimes necessary) for this early media to use the candidate pair ultimately selected by ICE connectivity checks. For this reason, ICE provides an early media mechanism that allows for a candidate pair to be used in one direction prior to its promotion to active in a subsequent offer/answer exchange. Note that, with ICE, early media pertains to media sent to a candidate pair until its promotion to active in a subsequent offer/answer exchange. This is a broader definition than is used in [29], which defines early media as media sent prior to acceptance of a call.

As a consequence of the connectivity checks, an agent will change the states for each transport address pair, and consequently, for the candidate pairs. When a candidate pair becomes Valid or Recv-Valid, and the candidate pair is not equal to the active candidate pair, and the agent is in the role of answerer for that candidate pair, the agent checks the position of that pair in the candidate pair priority ordered list. If it is the first, the agent selects this candidate pair for early media. If this candidate pair is not the first on the candidate pair priority ordered list, but is higher priority than the active candidate pair, and the early media wait-state timer has not

yet been set, the agent sets this timer to Tws seconds. Tws SHOULD be configurable, and SHOULD have a default of 100ms. This timer allows for a higher priority connectivity check to complete, in the event its STUN Binding Request or Response was lost or delayed in the network. If, prior to the wait-state timer firing, another connectivity check completes and a candidate pair enters the Valid or Recv-Valid states, there is no need to reset or cancel the timer. Once the timer fires, the agent SHOULD select the highest priority candidate pair in the Valid or Recv-Valid state for which the agent has the role of answerer, and use that candidate pair for early media.

ICE processing will ensure that, under almost all circumstances, the candidate pair selected by the answerer for early media will also be the one selected by the offerer for eventual promotion to active. The early media state implies that the answerer knows that this candidate pair is to be used, but the offerer doesn't know yet that it will eventually be validated. It is for this reason that the candidate pair can be used for early media.

If a candidate pair is selected for early media, an agent MAY send media on that candidate pair, even if it is not the same as the active candidate pair. However, to deal with cases in which the offerer and answerer do not agree on the eventual selection of this candidate for promotion to active (a rare but possible case), the agent MUST discontinue using the candidate pair for sending media Tlo seconds after the answer has been reliably delivered. An answer is considered reliably delivered when the agent receives a confirmation that is has been delivered. In the case of an answer delivered in a 200 OK to an offer in an INVITE (in the SIP case), the answer is considered reliably delivered upon receipt of the ACK. Tlo SHOULD be configurable and SHOULD have a default of 5 seconds. This time represents the amount of time it should take the offerer to perform its connectivity checks, arrive at the same conclusion about the viability of the early candidate, and then generate an updated offer promoting it to active. If, after Tlo seconds, no updated offer arrives, the answerer MUST cease using the early candidate. Media MAY be sent to the active candidate pair if it is in the Valid or Recv-Valid state.

If an updated offer does arrive prior to the expiration of the timer, the agent MUST execute the procedures in Section 7.11.2, which will result in the selection of a candidate for the m/c-line in the answer. At that point, the procedures of this section SHOULD be restarted by the answerer. This implies that the active candidate pair, if Valid or Recv-Valid, will be used. If a higher priority candidate pair subsequently enters the Valid or Recv-Valid state, it may end up being used as an early candidate.

To use a candidate pair, whether it is early or active, media is sent to the IP addresses and ports of the components in the remote candidate, and sends that media from the IP addresses and ports of the components in the native candidate. Transport addresses are paired up based on component ID. For example, if a remote candidate has two components R1 and R2, and the native candidate has two components L1 and L2, media packets are sent from L1 to R1 and from L2 to R2. This provides a property known as symmetry. This symmetric behavior MUST be followed by an agent even if its peer in the session doesn't support ICE.

The definition of sending media "from" a particular transport address depends on the type of transport address. In the case of a server reflexive transport address, this means that the RTP packets are sent from the local transport address used to obtain the STUN address. In the case of a relayed transport address, this means that media packets are sent through the relay server (for STUN relays, this would be using the Send request). For local transport addresses, media is sent from that local transport address. For peer reflexive transport addresses, media is sent from the local transport address used to obtain the reflexive address.

ICE has interactions with jitter buffer adaptation mechanisms. An RTP stream can begin using one candidate, and switch to another one. The newer candidate may result in RTP packets taking a different path through the network - one with different delay characteristics. As discussed below, agents are encouraged to re-adjust jitter buffers when there are changes in source or destination address. Furthermore, many audio codecs use the marker bit to signal the beginning of a talkspurt, for the purposes of jitter buffer adaptation. For such codecs, it is RECOMMENDED that the sender change the marker bit when an agent switches transmission of media from one candidate pair to another.

7.14 Receiving Media

ICE implementations MUST be prepared to receive media on a candidate pair if it is in the role of offerer for that candidate pair, even if that candidate pair is not currently active. This is a consequence of the early media mechanism described in the previous section.

If an agent determines that its peer supports ICE (an offerer knows this when the answer contains a=candidate attributes), it SHOULD discard any media packets received on a candidate pair prior to the candidate pair entering the Send Valid state. This helps eliminate certain attacks, as discussed in Section 13.

It is RECOMMENDED that, when an agent receives an RTP packet with a

new source or destination IP address for a particular media stream, that the agent re-adjust its jitter buffers.

RFC 3550 [23] describes an algorithm in Section 8.2 for detecting SSRC collisions and loops. These algorithms are based, in part, on seeing different source IP addresses and ports with the same SSRC. However, when ICE is used, such changes will naturally occur as the media streams switch between candidates. An agent will be able to determine that a media stream is from the same peer as a consequence of the STUN exchange that proceeds media transmission. Thus, if there is a change in source IP address and port, but the media packets come from the same peer agent, this SHOULD NOT be treated as an SSRC collision.

8. Guidelines for Usage with SIP

SIP [2] makes use of the offer/answer model, and is one of the primary targets for usage of ICE. SIP allows for offer/answer exchanges to occur in many different combinations of messages, including INVITE/200 OK and 200 OK/ACK. When support for reliable provisional responses (RFC 3262 [11]) and UPDATE (RFC 3311 [27]) are added, additional combinations of messages that can be used for offer/answer exchanges are added. As such, this section provides some guidance on good ways to make use of SIP with ICE.

ICE requires a series of STUN-based connectivity checks to take place between endpoints. These checks start from the answerer on generation of its answer, and start from the offerer when it receives the answer. These checks can take time to complete, and as such, the selection of messages to use with offers and answers can effect perceived user latency. Two latency of figures are of particular interest. These are the post-pickup delay and the post-dial delay. The post-pickup delay refers to the time between when a user "answers the phone" and when any speech they utter can be delivered to the caller. The post-dial delay refers to the time between when a user enters the destination address for the user, and ringback begins as a consequence of having succesfully started ringing the phone of the called party.

To reduce post-dial delays, it is RECOMMENDED that the caller begin gathering candidates prior to actually sending its initial INVITE. This can be started upon user interface cues that a call is pending, such as activity on a keypad or the phone going offhook.

To reduce post-pickup delays, ICE allows for media to be sent from the answerer to the offerer on a candidate pair, prior to its promotion to active. However, this requires the answerer to have generated its answer and sent it. In most cases, it will require

this answer to be received by the offerer. The reason is that connectivity checks or RTP packets from the answerer to the offerer will not be forwarded by NATs towards the offerer until the offerer has established a permission in the NAT by generating a packet towards the answerer.

For this reason, if an offer is received in an INVITE request, the UAS SHOULD immediately gather its candidates and then generate an answer in a provisional response. When reliable provisional responses are not used, the SDP in the provisional response is not formally the answer; the value in the 200 OK is the actual answer. However, RFC 3261 allows for SDP to appear in an unreliable provisional response, in which case its value has to be identical to the value placed in the 200 OK. Thus, we refer to the SDP in the provisional response, even when unreliable, as the answer. To deal with possible losses of the provisional response, it SHOULD be retransmitted until some indication of receipt. This indication can either be through PRACK [11], or through the receipt of a STUN Binding Request with a correct username and password. Even if PRACK is not used, the provisional response SHOULD be retransmitted using the exponential backoff described in [11]. Furthermore, once the answer has been sent, the agent SHOULD begin its connectivity checks. Once a candidate reaches the Valid or Recv-Valid state, the UAS has a known-valid path for media packets towards the UAC. This point is called the connected point in ICE.

Once the UAS reaches the connected point, media can be sent from the UAS towards the UAC without any additional delays. However, between the receipt of the INVITE and the connected point, any media that needs to be sent towards the caller (such as SIP early media [29] cannot be transmitted. For this reason, implementations MAY choose to delay alerting the called party until the connected point is reached. In the case of a PSTN gateway, this would mean that the setup message into the PSTN is delayed until the connected point. Doing this increases the post-dial delay, but has the effect of eliminating 'ghost rings'. Ghost rings are cases where the called party hears the phone ring, picks up, but hears nothing and cannot be heard. This technique works without requiring support for, or usage of, preconditions [7], since its a localized decision. It also has the benefit of guaranteeing that not a single packet of early media will get clipped. If an agent chooses to delay local alerting in this way, it SHOULD generate a 180 response once alerting begins.

A slight variation of this approach is to wait for a connectivity check to succeed to a higher priority candidate pair than the active one. This allows for the agent to only ever send media, early or otherwise, to a single candidate, which will work better with jitter buffers, at the expense of even greater post-dial delays.

Note that, prior to the promotion of a candidate pair to active, the offerer will not be able to send using the candidate pair. When used with SIP, if the initial offer is sent in the INVITE, and the answer is sent in both the provisional and final 200 OK response, the offerer will not be able to send media until it sends a re-INVITE and receives the 200 OK response to that re-INVITE. This can take several hundred milliseconds. If this latency is an issue (it is generally not considered an issue for voice systems), reliable provisional responses [11] MAY be used, in which case an UPDATE [27] can be used to send an updated offer prior to the call being answered.

As discussed in Section 13, offer/answer exchanges SHOULD be secured against eavesdropping and man-in-the-middle attacks. To do that, the usage of SIPS [2] is RECOMMENDED when used in concert with ICE.

9. Interactions with Forking

SIP allows INVITE requests carrying offers to fork, which means that they are delivered to multiple user agents. Each of those user agents then provides an answer to the offer in the INVITE. The result is that a single offer generated by the UAC produces multiple answers.

ICE interacts very well with forking. Indeed, ICE fixes some of the problems associated with forking. Once the offer/answer exchange has completed, the UAC will have an answer from each UAS that received the INVITE. The ICE connectivity checks that ensue will carry transport address pair IDs that correlate each of those checks (and thus their corresponding IP addresses and ports) with a specific remote user agent. As these checks happen before any media is transmitted, ICE allows a UAC to disambiguate subsequent media traffic by looking at the source IP address and port, and then correlate that traffic with a particular remote UA. When SIP is used without ICE, the incoming media traffic cannot be disambiguated without an additional offer/answer exchange.

10. Interactions with Preconditions

Because ICE involves multiple addresses and pre-session activities, its interactions with preconditions merits further discussion.

Quality of Service (QoS) preconditions, which are defined in RFC 3312 [7] and RFC 4032 [8], apply only to the IP addresses and ports listed in the m/c lines in an offer/answer. If ICE changes the address and port where media is received, this change is reflected in the m/c lines of a new offer/answer. As such, it appears like any other re-INVITE would, and is fully treated in RFC 3312 and 4032, which

applies without regard to the fact that the m/c lines are changing due to ICE negotiations ocurring "in the background".

However, usage of early candidates with QoS preconditions is NOT RECOMMENDED, since QoS will only be reserved for the candidate pair in the m/c-line. An agent SHOULD only send to the active candidate (once it enters the Valid or Recv-Valid states) if QoS preconditions are used for a media session.

ICE also has (purposeful) interactions with connectivity preconditions [30]. Those interactions are described there.

11. Examples

This section provides two examples. One is a very basic example, and the other is more elaborate. A common configuration and setup is used in both cases.

Two agents, L and R, are using ICE. Both agents have a single IPv4 interface, and are configured with a single STUN server each (indeed, the same one for each). This STUN server supports both the Binding Discovery usage and the Relay usage. Agent L is behind a NAT, and agent R is on the public Internet.

To facilitate understanding, transport addresses are listed in a mnemonic form. This form is entity-type-seqno, where entity refers to the entity whose interface the transport address is on, and is one of "L", "R", "STUN", or "NAT". The type is either "PUB" for transport addresses that are public, and "PRIV" for transport addresses that are private. Finally, seq-no is a sequence number that is different for each transport address of the same type on a particular entity.

The STUN server has advertised transport address STUN-PUB-1 for both the binding discovery usage and the relay usage.

In addition, candidate IDs are also listed in mnemonic form. Agent L uses candidate ID L1 for its local candidate, L2 for its server reflexive candidate, and L3 for its relayed candidate. Agent R uses R1 for its local candidate and R2 for its relayed candidate. The password is LPASS for each candidate from agent L, and RPASS for each candidate from agent R.

In example SDP messages, $TADDR.IP is used to refer to the value of the IP address of the transport address with mnemonic name "taddr". Similarly, $TADDR.PORT is used to refer to the value of the port of the transport address with mnemonic name "TADDR".

In the call flow itself, STUN messages are annotated with several attributes. The "S=" attribute indicates the source transport address of the message. The "D=" attribute indicates the destination transport address of the message. The "MA=" attribute is used in STUN Binding Response messages, STUN Binding Response messages carried in a STUN Send Request or Data Indication, and in a Allocate Response, and refers to the reflexive transport address derived from the XOR-MAPPED-ADDRESS attribute. The "RA=" attribute is used in STUN Data Indications, and refers to the value of the REMOTE-ADDRESS attribute. The "U=" attribute is used in STUN Requests, and corresponds to the STUN USERNAME. The "DA=" attribute is used in STUN Send requests, and refers to the value of the DESTINATION-ADDRESS attribute. The "R=" attribute is used in Allocate responses, and it indicates the value of the RELAY-ADDRESS attribute.

The call flow examples omit STUN authentication operations.

11.1 Basic Example

In this example, the NAT has the address and port independent mapping property and the address dependent permission property. Neither agent is using the STUN relay usage, only the binding discovery usage. As a consequence, agent L will end up with two candidates - a local candidate and a server reflexive candidate. Agent R will have one - a local candidate (the reflexive candidate will be identical to the local one, and thus discarded). The agents are seeking to communicate using a single RTP-based voice stream. RTCP is not used. As a consequence, each candidate has one component.

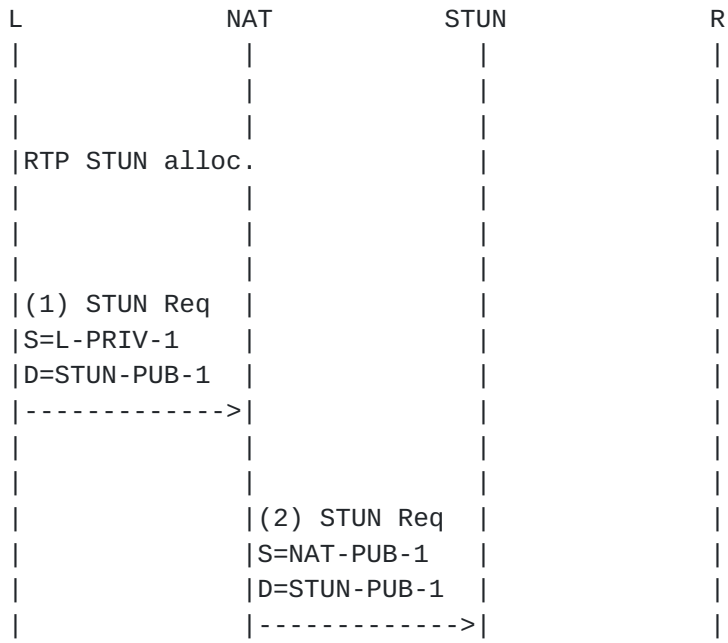

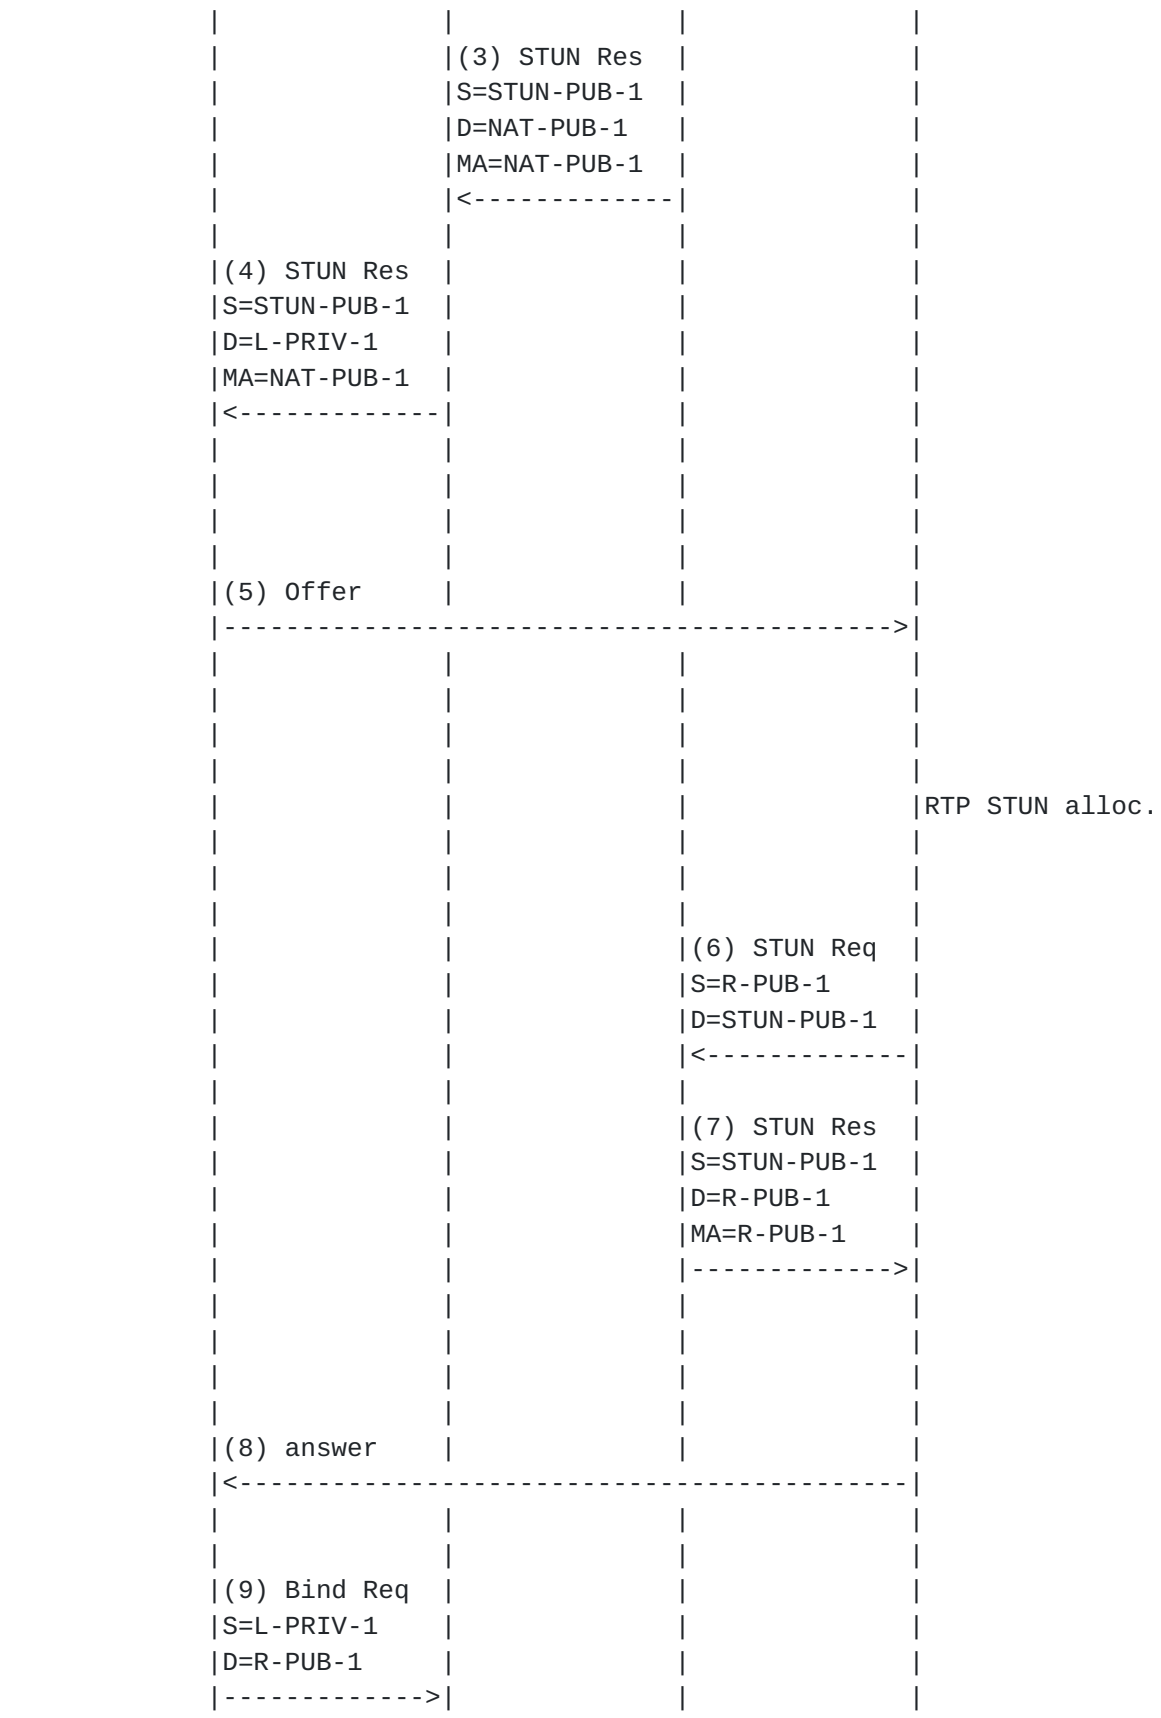
(3) STUN Res
S=STUN-PUB-1
D=NAT-PUB-1
MA=NAT-PUB-1
(4) STUN Res
S=STUN-PUB-1
D=L-PRIV-1
MA=NAT-PUB-1
(5) Offer
RTP STUN alloc.
(6) STUN Req
S=R-PUB-1
D=STUN-PUB-1
(7) STUN Res
S=STUN-PUB-1
D=R-PUB-1
MA=R-PUB-1
(8) answer
(9) Bind Req
S=L-PRIV-1
D=R-PUB-1

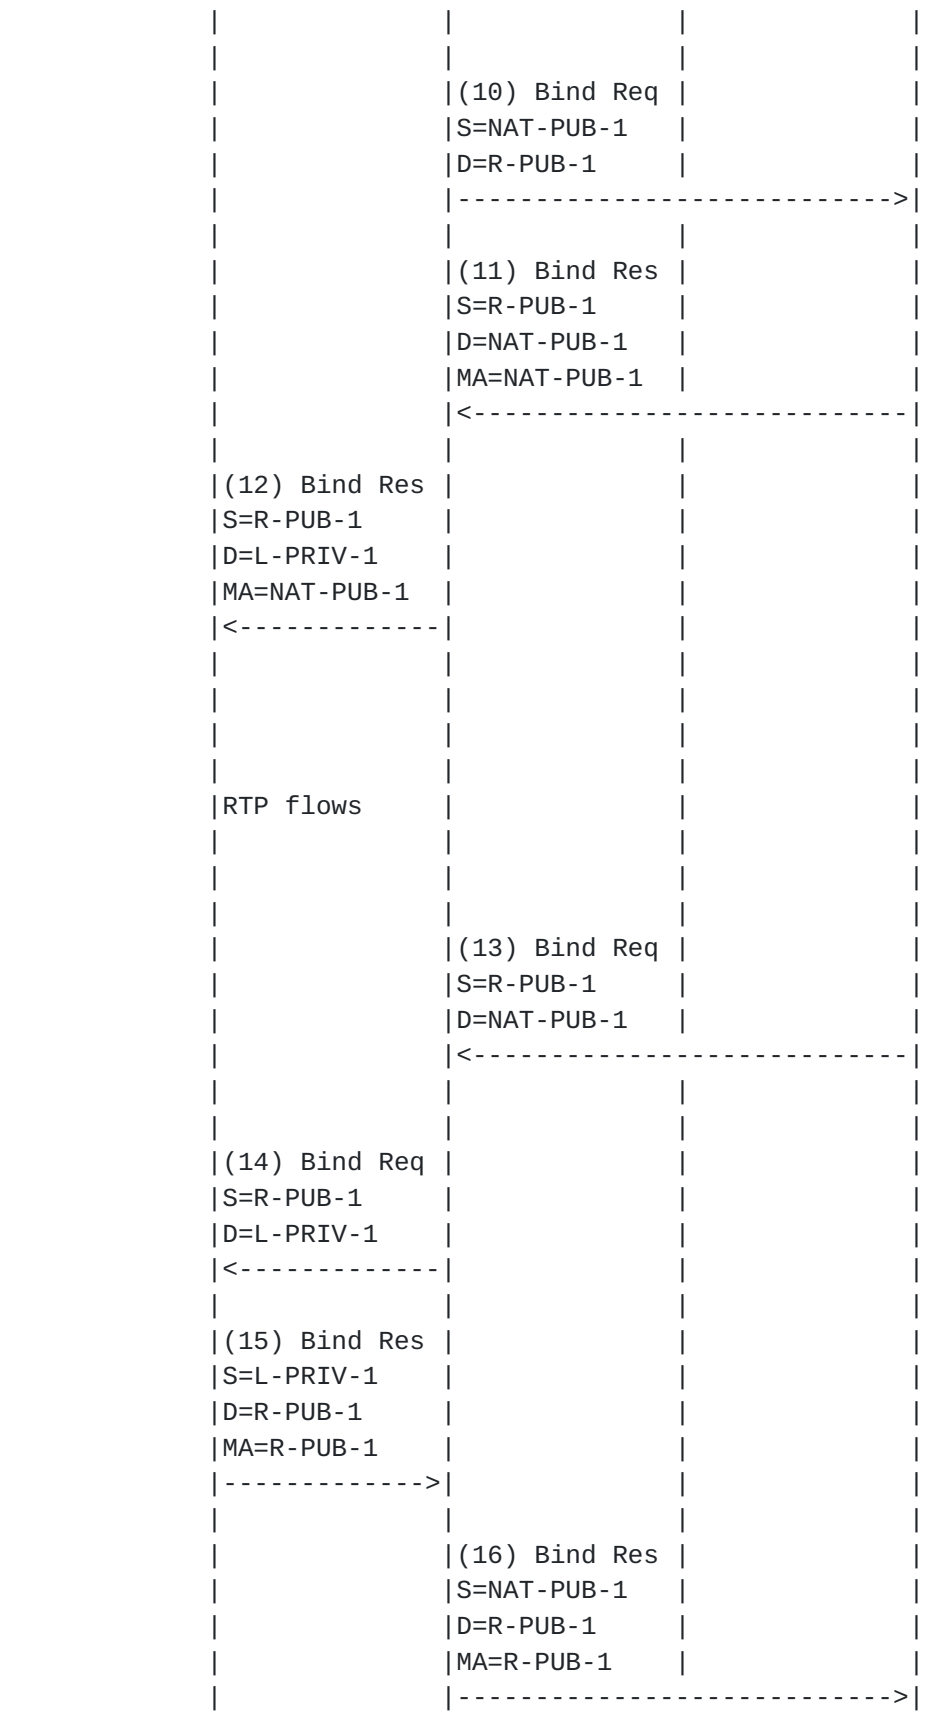
(10) Bind Req
S=NAT-PUB-1
D=R-PUB-1
(11) Bind Res
S=R-PUB-1
D=NAT-PUB-1
MA=NAT-PUB-1
(12) Bind Res
S=R-PUB-1
D=L-PRIV-1
MA=NAT-PUB-1
RTP flows
(13) Bind Req
S=R-PUB-1
D=NAT-PUB-1
(14) Bind Req
S=R-PUB-1
D=L-PRIV-1
(15) Bind Res
S=L-PRIV-1
D=R-PUB-1
MA=R-PUB-1
(16) Bind Res
S=NAT-PUB-1
D=R-PUB-1
MA=R-PUB-1

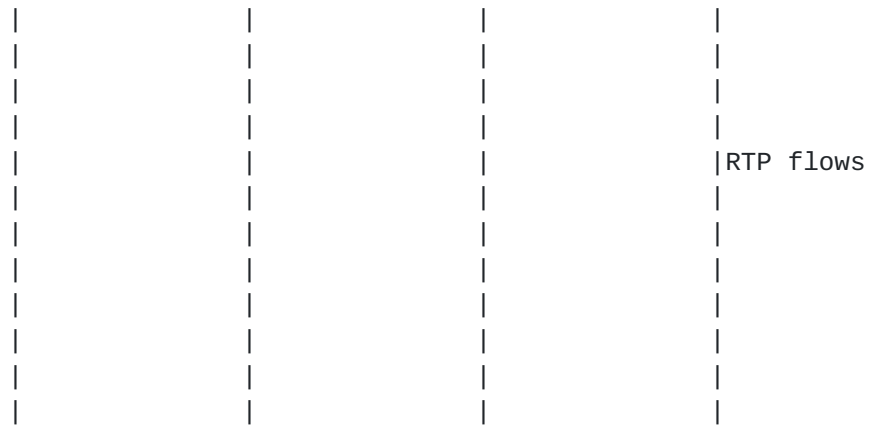

Figure 8

First, agent L obtains a server reflexive transport address for its RTP packets (messages 1-4). Recall that the NAT has the address and port independent mapping property. Here, it creates a binding of NAT-PUB-1 for this UDP request, and this becomes the server reflexive transport address for RTP, the sole component of its server reflexive candidate.

With its two candidates, agent L prioritizes them, choosing the local candidate as highest priority, followed by the server reflexive candidate. It chooses its server reflexive candidate as the active candidate, and encodes it into the m/c-line. The resulting offer (message 5) looks like:

```
v=0
o=jdoe 2890844526 2890842807 IN IP4 $L-PRIV-1.IP
s=
c=IN IP4 $STUN-PUB-1.IP
t=0 0
a=ice-pwd:$LPASS
m=audio $STUN-PUB-1.PORT RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=candidate $L1 1 UDP 1.0 $L-PRIV-1.IP $L-PRIV-1.PORT
a=candidate $L2 1 UDP 0.7 $NAT-PUB-1.IP $NAT-PUB-1.PORT
```

This offer is received at agent R. Agent R will gather its server reflexive transport address (messages 6-7). Since R is not behind a NAT, this address is identical to its local transport address, and thus does not represent a separate candidate. It therefore ends up with a single local candidate with a single component for RTP. Its resulting answer looks like:

```
v=0
o=bob 2808844564 2808844564 IN IP4 $R-PUB-1.IP
s=
c=IN IP4 $R-PUB-1.IP
t=0 0
a=ice-pwd:$RPASS
m=audio $R-PUB-1.PORT RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=candidate $R1 1 UDP 1.0 $R-PUB-1.IP $R-PUB-1.PORT
```

Next, agents L and R form candidate pairs and the transport address check ordered list. This list will start with the single component in the currently active candidate pair, L2:1:R1:1. Agent L begins its connectivity checks (messages 9-12), which succeed, placing the transport address pair and resulting candidate pair into the Recv-Valid state. Media can now flow. When agent R receives this request (message 10), the state of the candidate pair moves to Send-Valid. Agent R begins its connectivity checks (messages 13-16). When the check arrives at the NAT (message 13), it is permitted to pass since a permission was created towards $R-PUB-1 as a consequence of message 10. This check arrives at agent L, which generates a success response (message 11), and updates the state of the candidate pair to Valid. This response arrives at agent R, which also updates the state of the candidate pair to valid. Now, media can flow from agent R to agent L as well.

11.2 Advanced Example

In this more advanced example, The NAT has address and port dependent mapping and filtering properties. Both agents use the STUN relay usage in addition to the binding discovery usage. As a consequence, agent L will end up with three candidates - a local candidate, a relayed candidate, and a server reflexive candidate. Agent R will have two - a local candidate and a relayed candidate (the server reflexive candidate will equal the local candidate and thus not be used). The agents are seeking to communicate using a single RTP-based voice stream, but are using RTCP. As a consequence, each candidate has two components - one for RTP and one for RTCP.

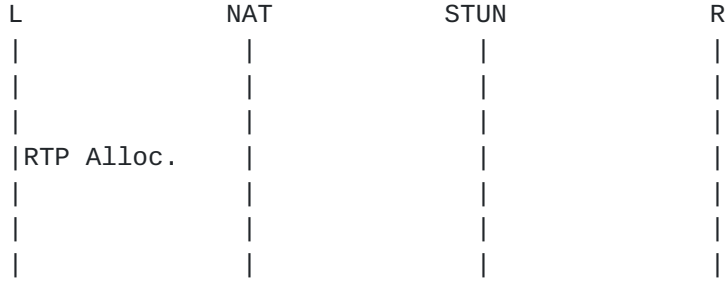

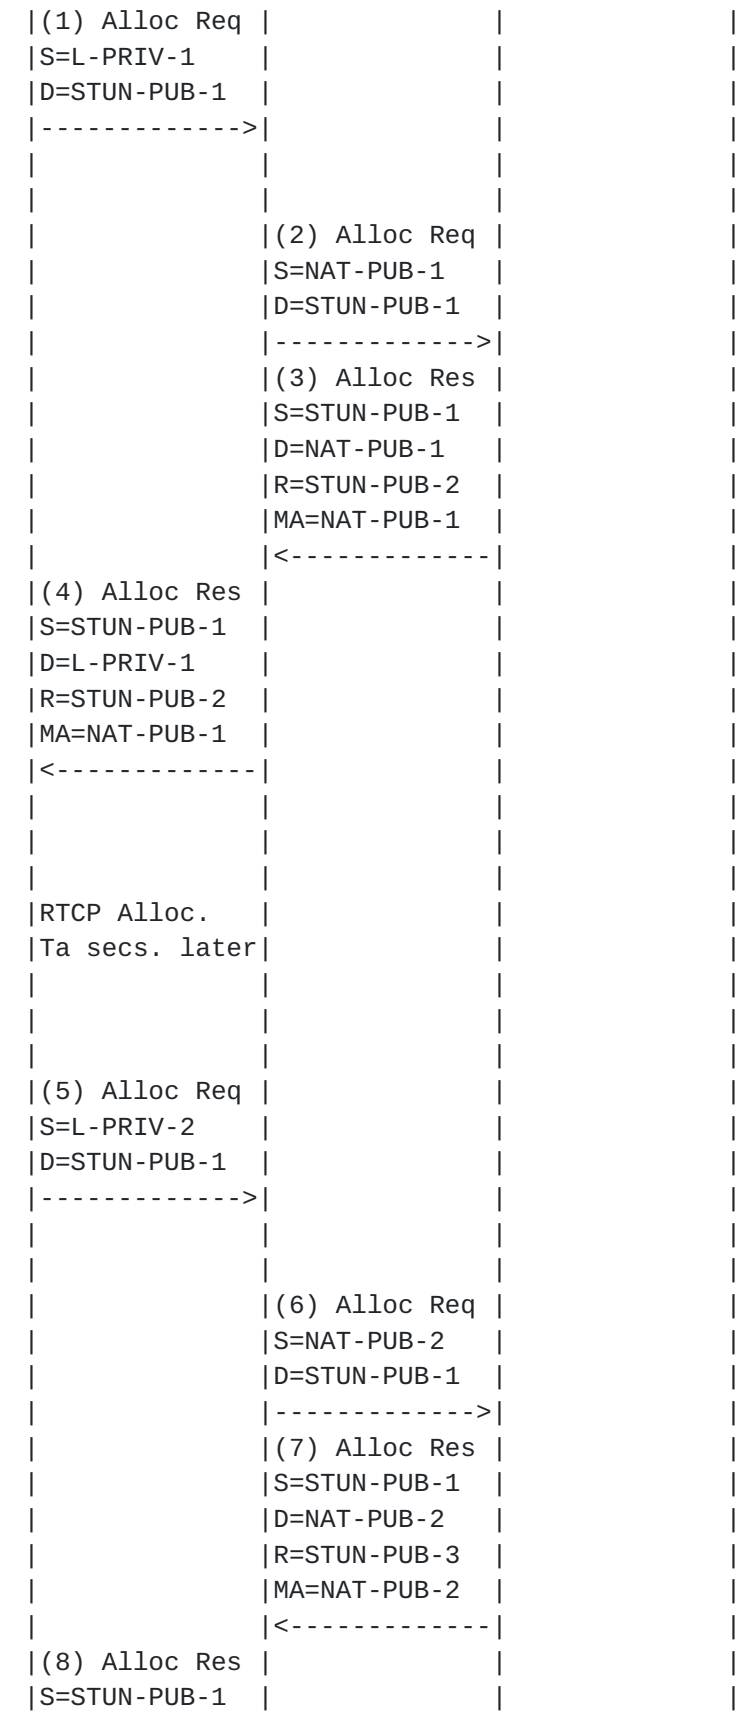
(1) Alloc Req
S=L-PRIV-1
D=STUN-PUB-1
(2) Alloc Req
S=NAT-PUB-1
D=STUN-PUB-1
(3) Alloc Res
S=STUN-PUB-1
D=NAT-PUB-1
R=STUN-PUB-2
MA=NAT-PUB-1
(4) Alloc Res
S=STUN-PUB-1
D=L-PRIV-1
R=STUN-PUB-2
MA=NAT-PUB-1
RTCP Alloc.
Ta secs. later
(5) Alloc Req
S=L-PRIV-2
D=STUN-PUB-1
(6) Alloc Req
S=NAT-PUB-2
D=STUN-PUB-1
(7) Alloc Res
S=STUN-PUB-1
D=NAT-PUB-2
R=STUN-PUB-3
MA=NAT-PUB-2
(8) Alloc Res
S=STUN-PUB-1

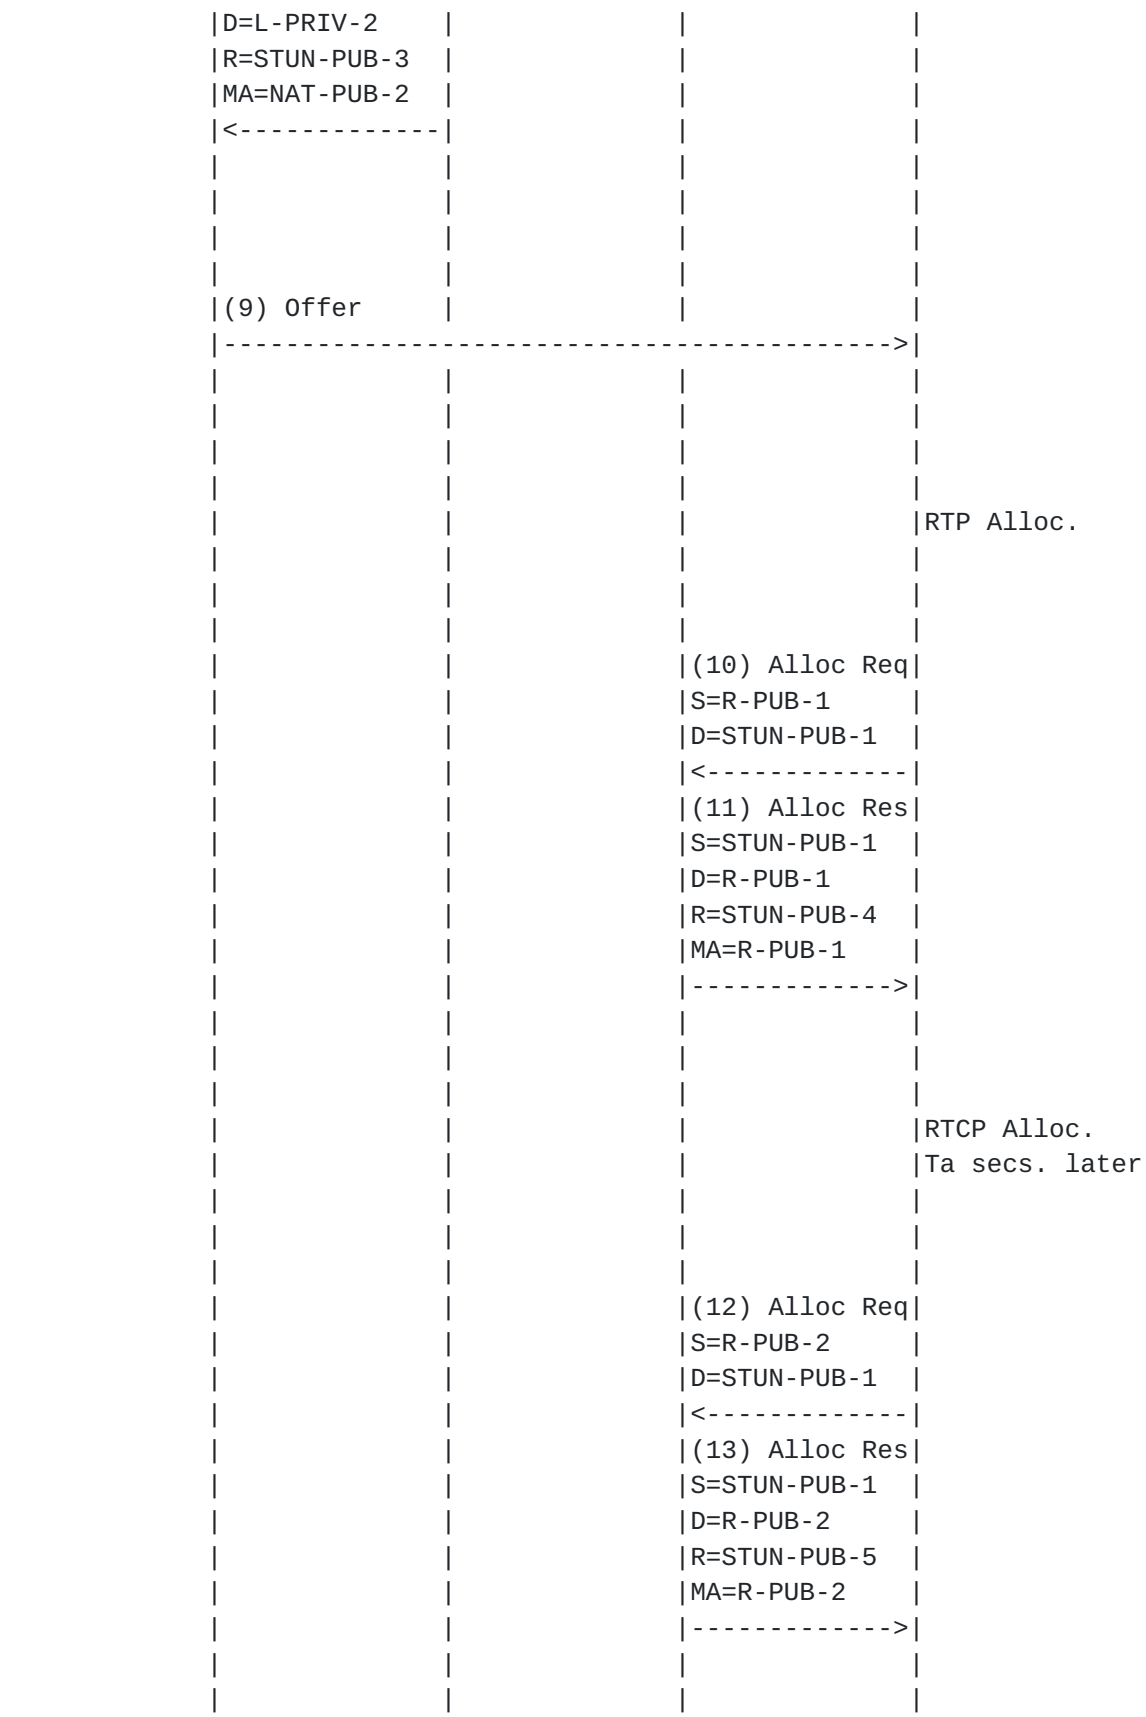
D=L-PRIV-2
R=STUN-PUB-3
MA=NAT-PUB-2
(9) Offer
RTP Alloc.
(10) Alloc Req
S=R-PUB-1
D=STUN-PUB-1
(11) Alloc Res
S=STUN-PUB-1
D=R-PUB-1
R=STUN-PUB-4
MA=R-PUB-1
RTCP Alloc.
Ta secs. later
(12) Alloc Req
S=R-PUB-2
D=STUN-PUB-1
(13) Alloc Res
S=STUN-PUB-1
D=R-PUB-2
R=STUN-PUB-5
MA=R-PUB-2

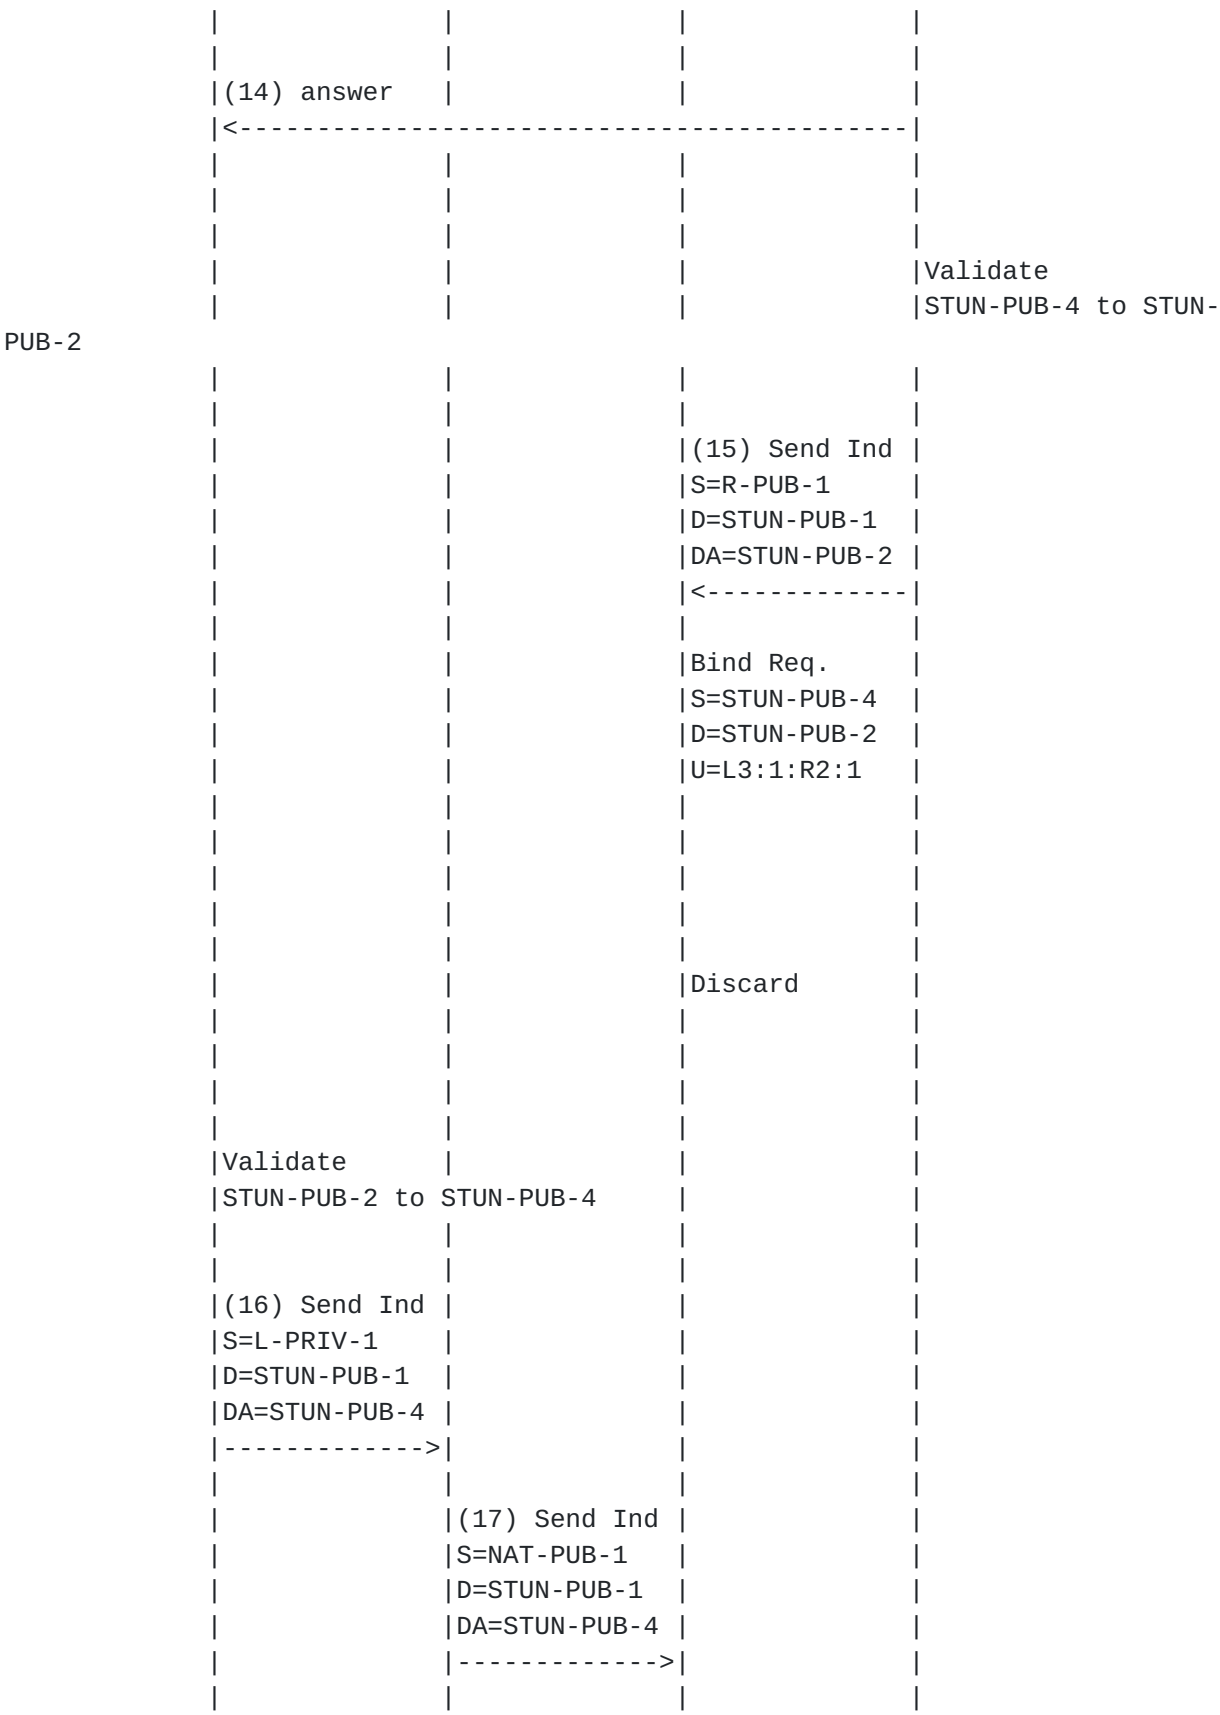
(14) answer
Validate
STUN-PUB-4 to STUN-PUB-2
(15) Send Ind
S=R-PUB-1
D=STUN-PUB-1
DA=STUN-PUB-2
Bind Req.
S=STUN-PUB-4
D=STUN-PUB-2
U=L3:1:R2:1
Discard
Validate
STUN-PUB-2 to STUN-PUB-4
(16) Send Ind
S=L-PRIV-1
D=STUN-PUB-1
DA=STUN-PUB-4
(17) Send Ind
S=NAT-PUB-1
D=STUN-PUB-1
DA=STUN-PUB-4

|Bind Req.

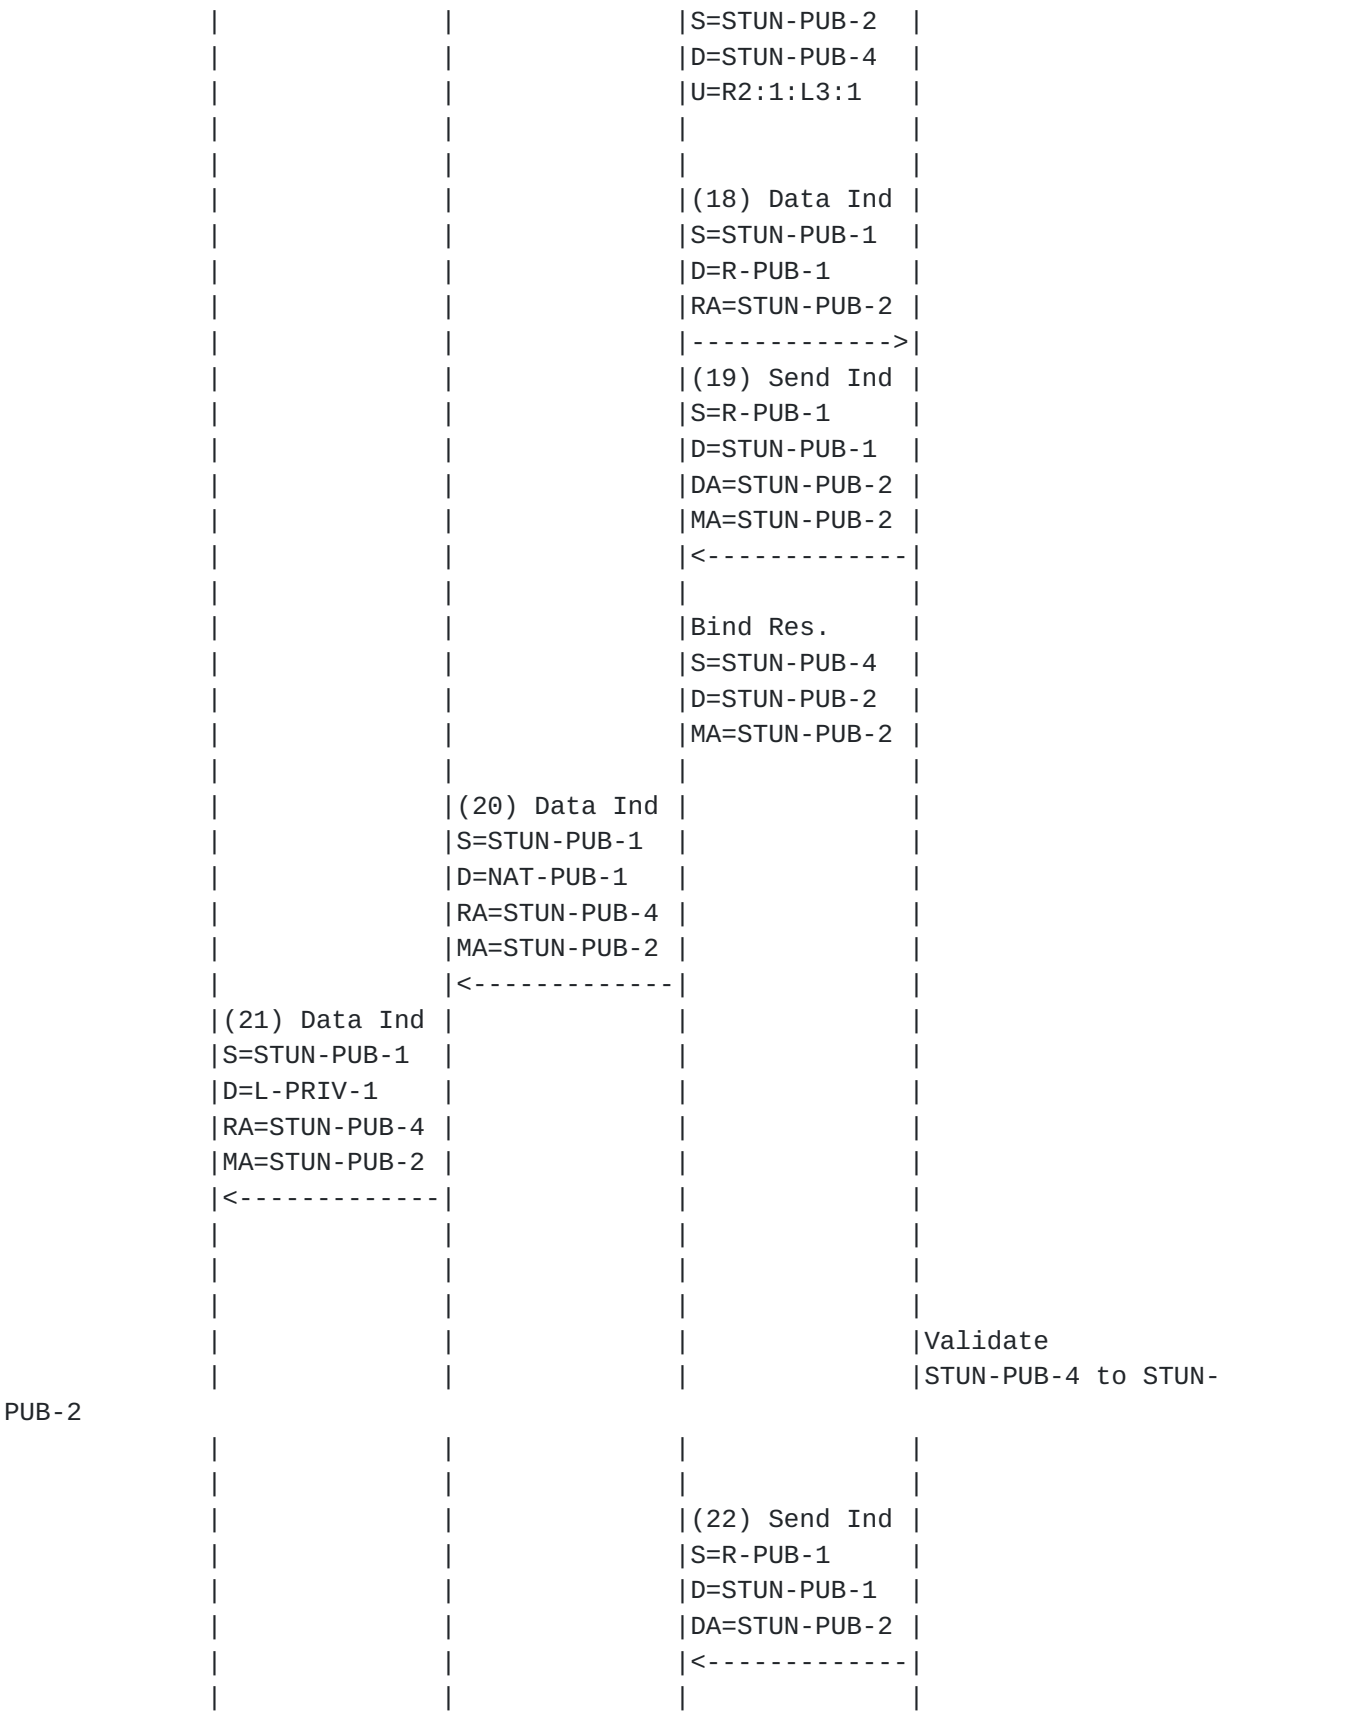
S=STUN-PUB-2
D=STUN-PUB-4
U=R2:1:L3:1
(18) Data Ind
S=STUN-PUB-1
D=R-PUB-1
RA=STUN-PUB-2
(19) Send Ind
S=R-PUB-1
D=STUN-PUB-1
DA=STUN-PUB-2
MA=STUN-PUB-2
Bind Res.
S=STUN-PUB-4
D=STUN-PUB-2
MA=STUN-PUB-2
(20) Data Ind
S=STUN-PUB-1
D=NAT-PUB-1
RA=STUN-PUB-4
MA=STUN-PUB-2
(21) Data Ind
S=STUN-PUB-1
D=L-PRIV-1
RA=STUN-PUB-4
MA=STUN-PUB-2
Validate
STUN-PUB-4 to STUN-PUB-2
(22) Send Ind
S=R-PUB-1
D=STUN-PUB-1
DA=STUN-PUB-2

|Bind Req.

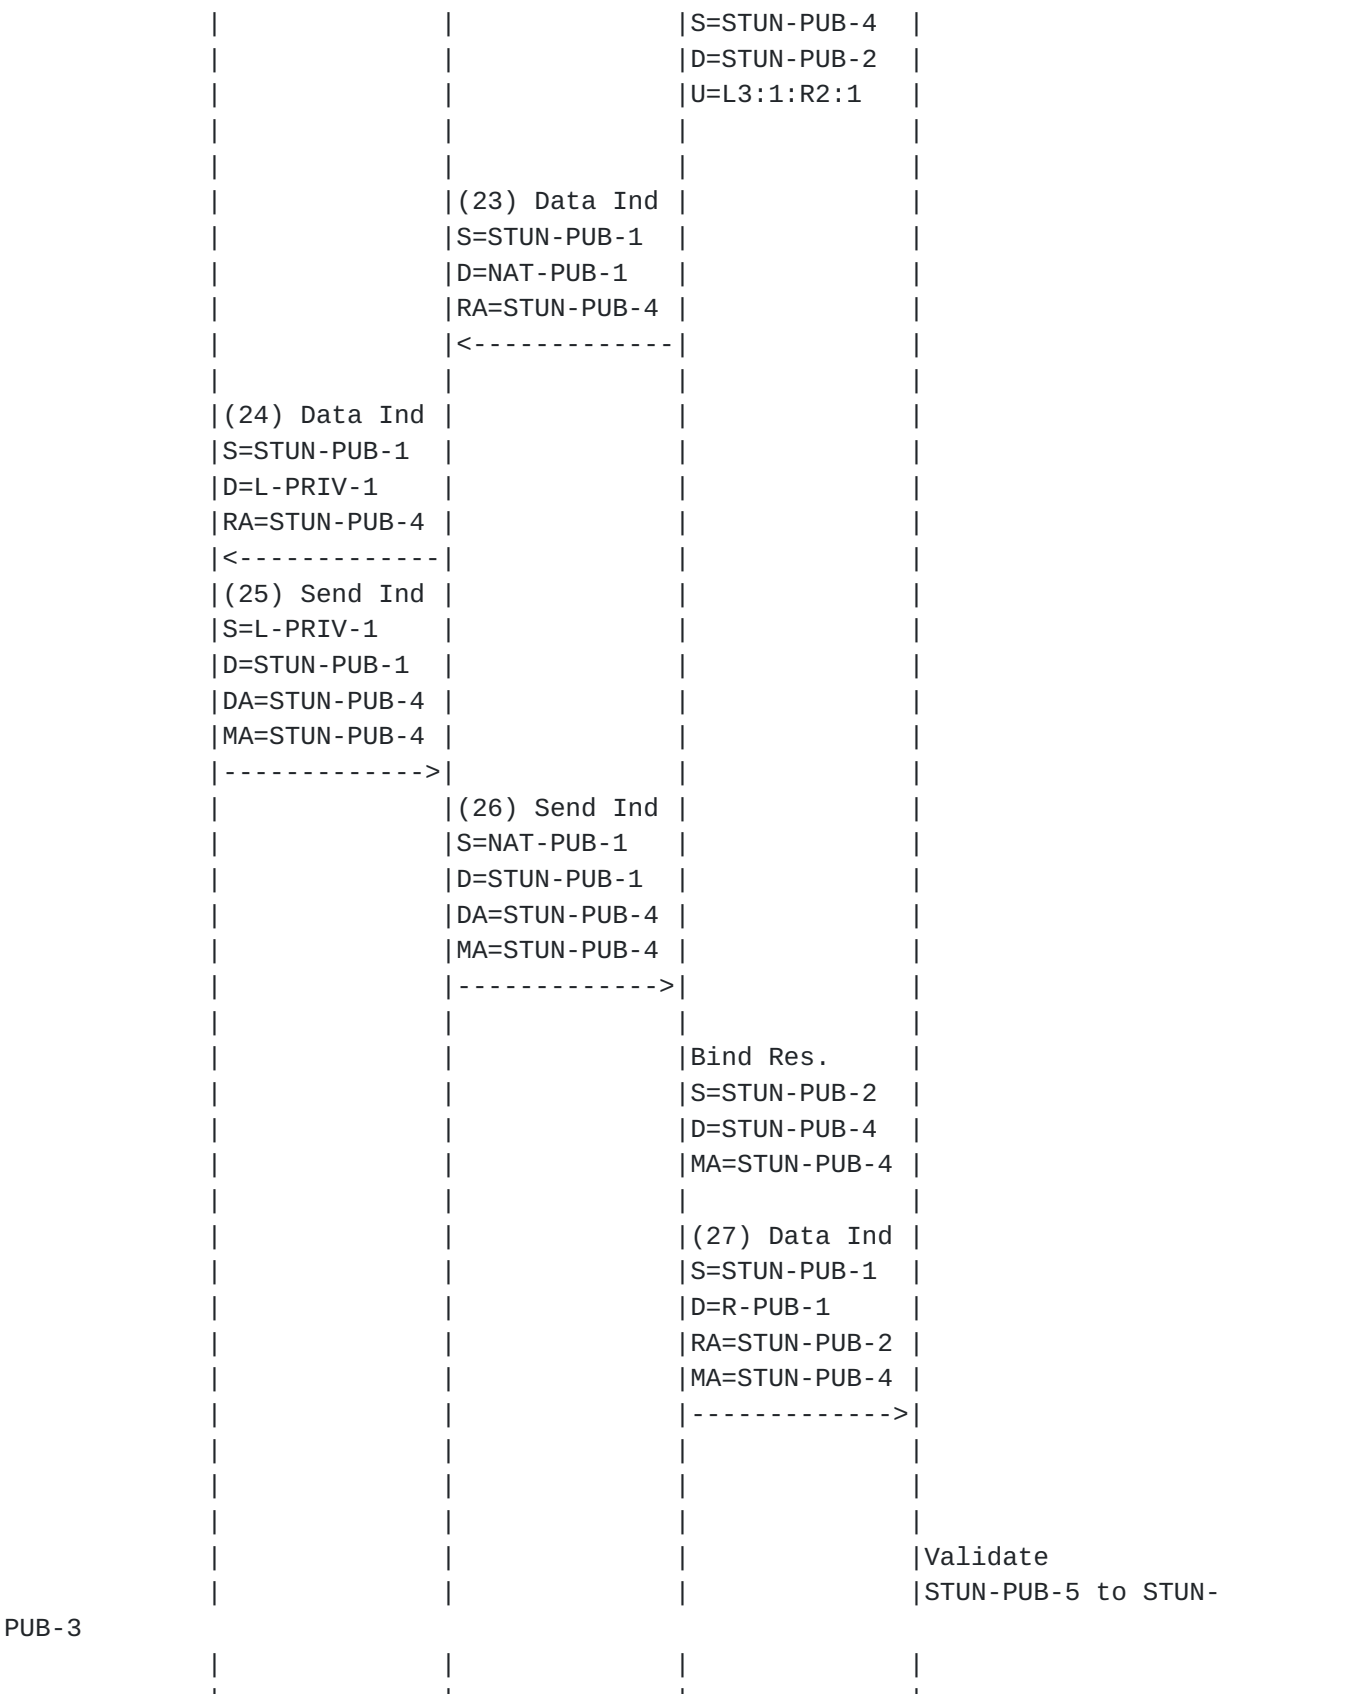
S=STUN-PUB-4
D=STUN-PUB-2
U=L3:1:R2:1
(23) Data Ind
S=STUN-PUB-1
D=NAT-PUB-1
RA=STUN-PUB-4
(24) Data Ind
S=STUN-PUB-1
D=L-PRIV-1
RA=STUN-PUB-4
(25) Send Ind
S=L-PRIV-1
D=STUN-PUB-1
DA=STUN-PUB-4
MA=STUN-PUB-4
(26) Send Ind
S=NAT-PUB-1
D=STUN-PUB-1
DA=STUN-PUB-4
MA=STUN-PUB-4
Bind Res.
S=STUN-PUB-2
D=STUN-PUB-4
MA=STUN-PUB-4
(27) Data Ind
S=STUN-PUB-1
D=R-PUB-1
RA=STUN-PUB-2
MA=STUN-PUB-4
Validate
STUN-PUB-5 to STUN-PUB-3

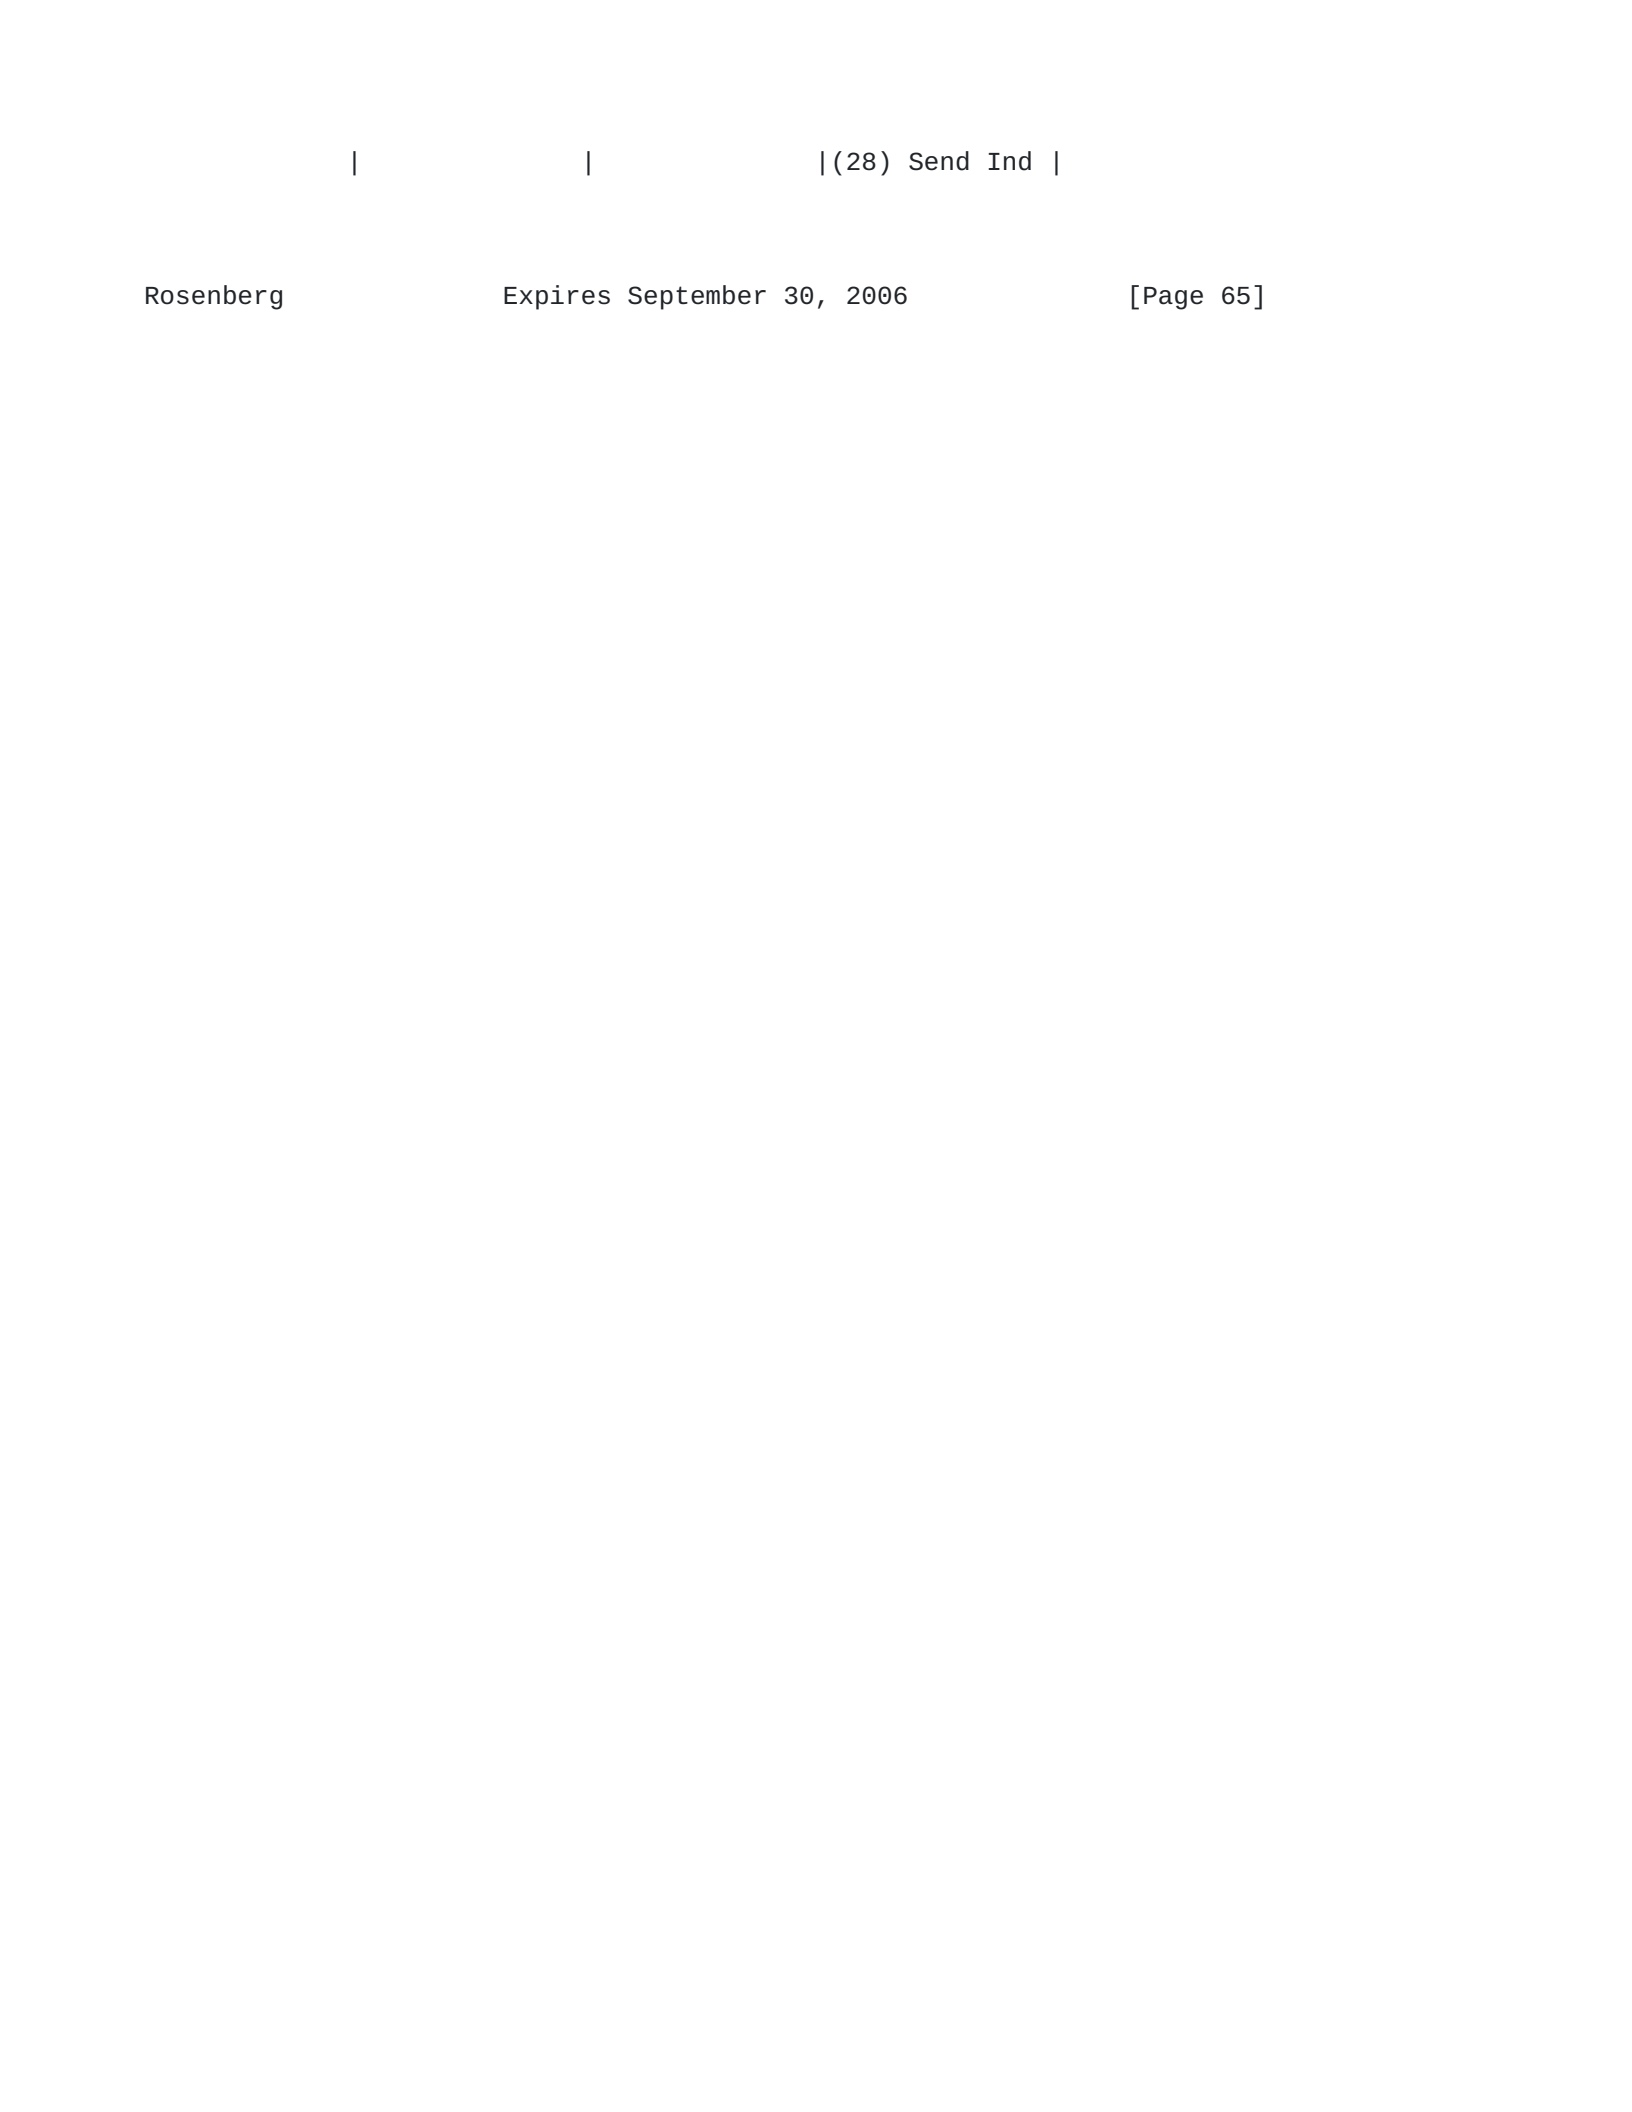
|(28) Send Ind |

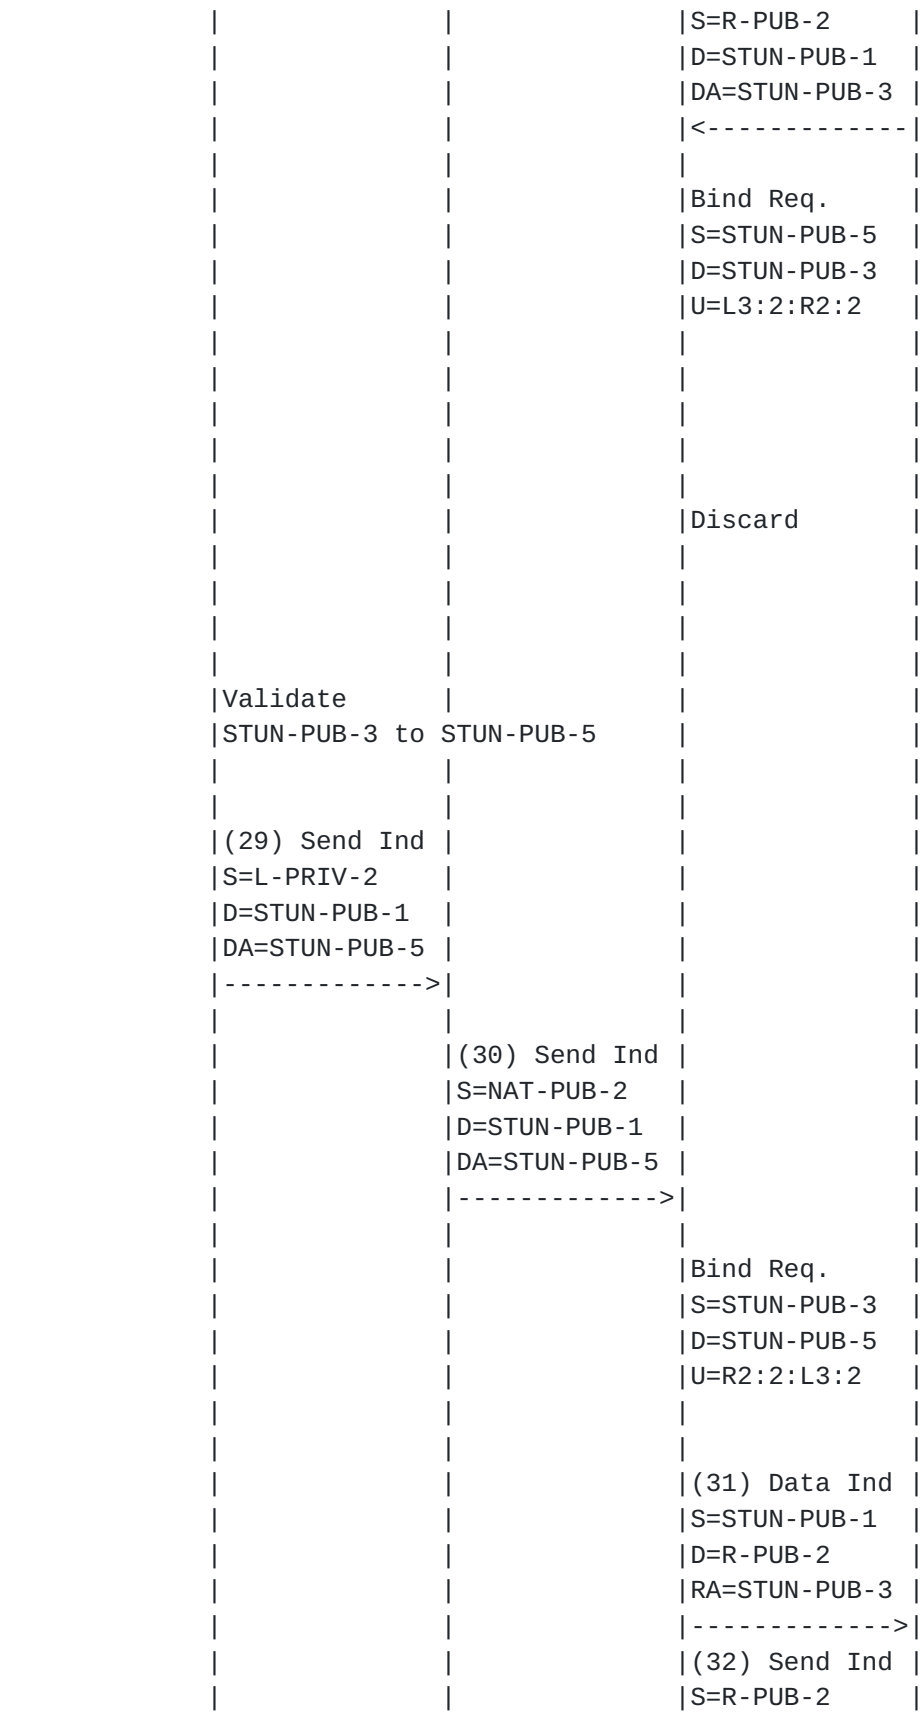
S=R-PUB-2
D=STUN-PUB-1
DA=STUN-PUB-3
Bind Req.
S=STUN-PUB-5
D=STUN-PUB-3
U=L3:2:R2:2
Discard
Validate
STUN-PUB-3 to STUN-PUB-5
(29) Send Ind
S=L-PRIV-2
D=STUN-PUB-1
DA=STUN-PUB-5
(30) Send Ind
S=NAT-PUB-2
D=STUN-PUB-1
DA=STUN-PUB-5
Bind Req.
S=STUN-PUB-3
D=STUN-PUB-5
U=R2:2:L3:2
(31) Data Ind
S=STUN-PUB-1
D=R-PUB-2
RA=STUN-PUB-3
(32) Send Ind
S=R-PUB-2

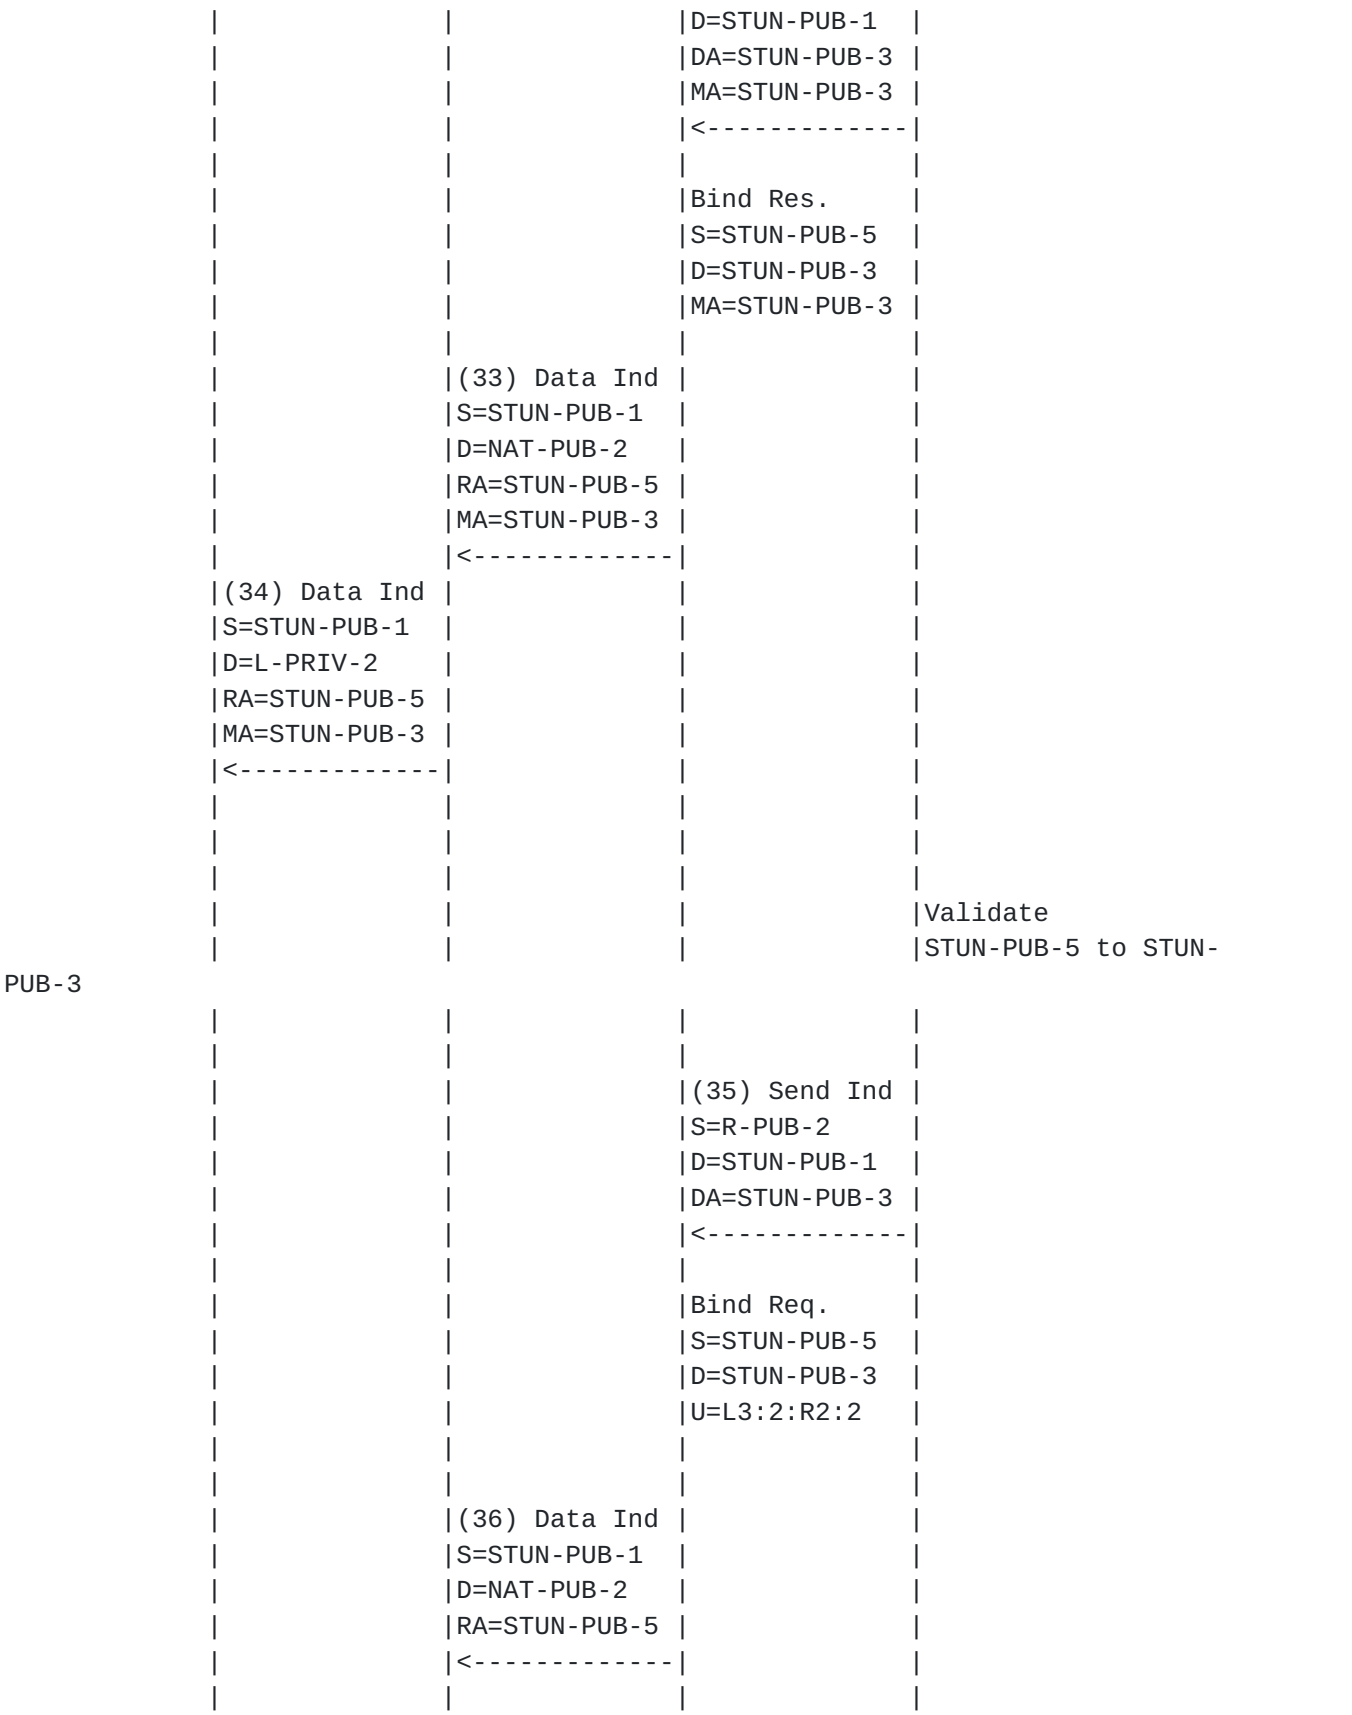
D=STUN-PUB-1
DA=STUN-PUB-3
MA=STUN-PUB-3
Bind Res.
S=STUN-PUB-5
D=STUN-PUB-3
MA=STUN-PUB-3
(33) Data Ind
S=STUN-PUB-1
D=NAT-PUB-2
RA=STUN-PUB-5
MA=STUN-PUB-3
(34) Data Ind
S=STUN-PUB-1
D=L-PRIV-2
RA=STUN-PUB-5
MA=STUN-PUB-3
Validate
STUN-PUB-5 to STUN-PUB-3
(35) Send Ind
S=R-PUB-2
D=STUN-PUB-1
DA=STUN-PUB-3
Bind Req.
S=STUN-PUB-5
D=STUN-PUB-3
U=L3:2:R2:2
(36) Data Ind
S=STUN-PUB-1
D=NAT-PUB-2
RA=STUN-PUB-5

|(37) Data Ind | | |

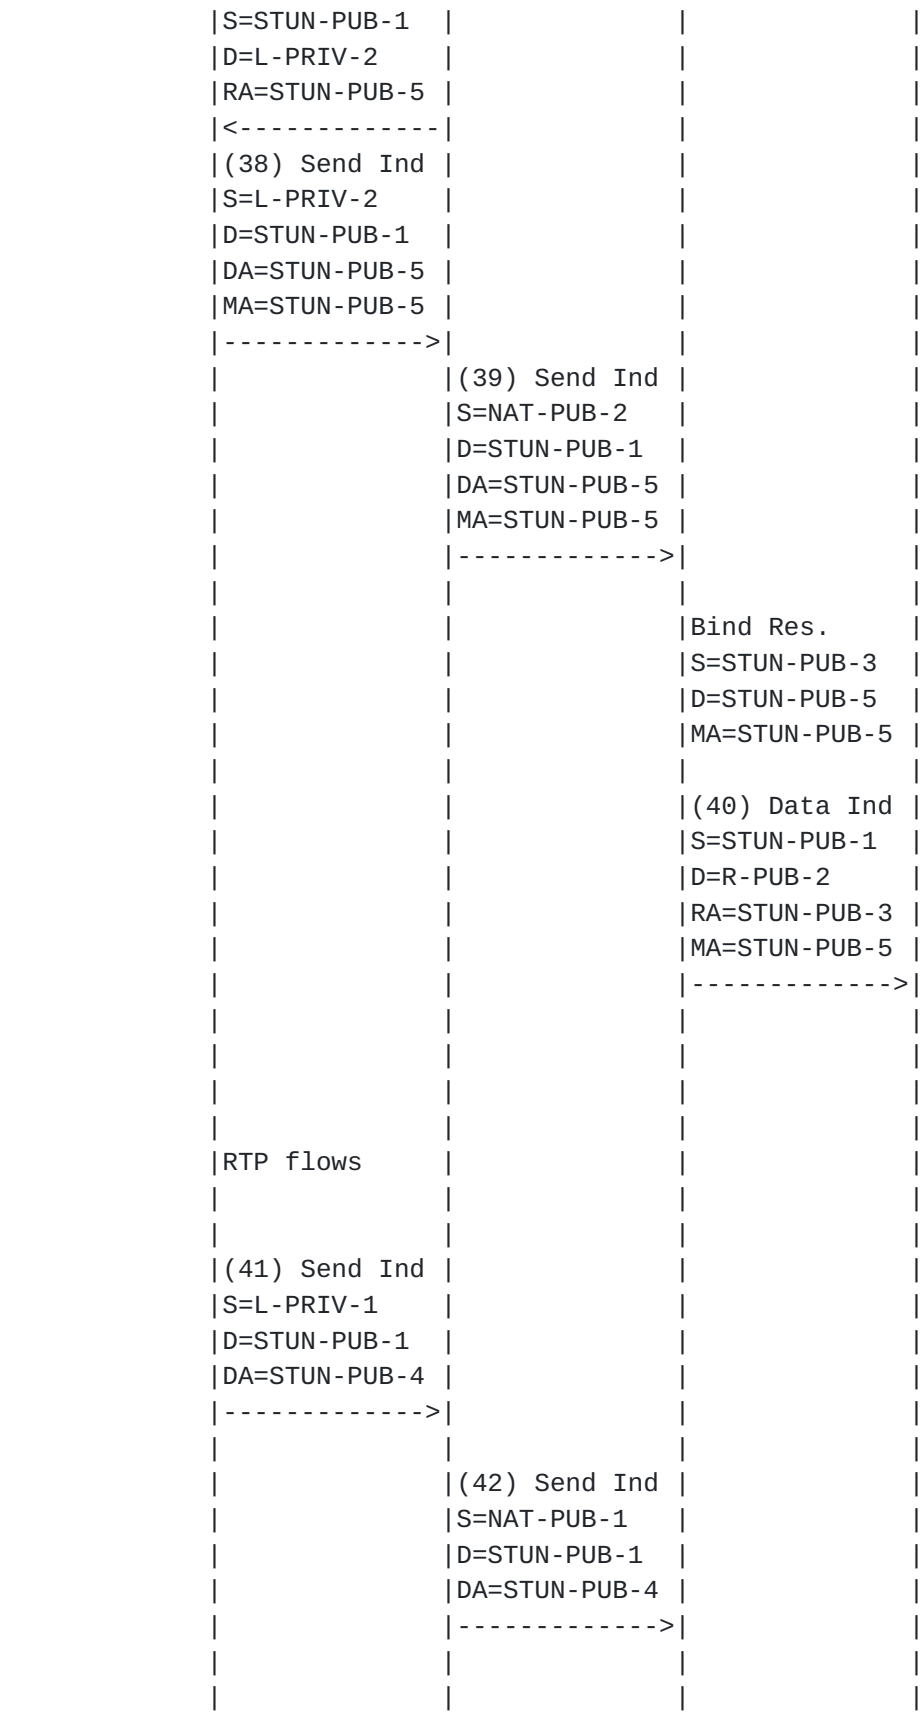
|S=STUN-PUB-1 |
|D=L-PRIV-2 |
|RA=STUN-PUB-5 |
|<----------|
|(38) Send Ind |
|S=L-PRIV-2 |
|D=STUN-PUB-1 |
|DA=STUN-PUB-5 |
|MA=STUN-PUB-5 |
|---------->|
|(39) Send Ind |
|S=NAT-PUB-2 |
|D=STUN-PUB-1 |
|DA=STUN-PUB-5 |
|MA=STUN-PUB-5 |
|---------->|
|Bind Res. |
|S=STUN-PUB-3 |
|D=STUN-PUB-5 |
|MA=STUN-PUB-5 |
|(40) Data Ind |
|S=STUN-PUB-1 |
|D=R-PUB-2 |
|RA=STUN-PUB-3 |
|MA=STUN-PUB-5 |
|---------->|
|RTP flows |
|(41) Send Ind |
|S=L-PRIV-1 |
|D=STUN-PUB-1 |
|DA=STUN-PUB-4 |
|---------->|
|(42) Send Ind |
|S=NAT-PUB-1 |
|D=STUN-PUB-1 |
|DA=STUN-PUB-4 |
|---------->|

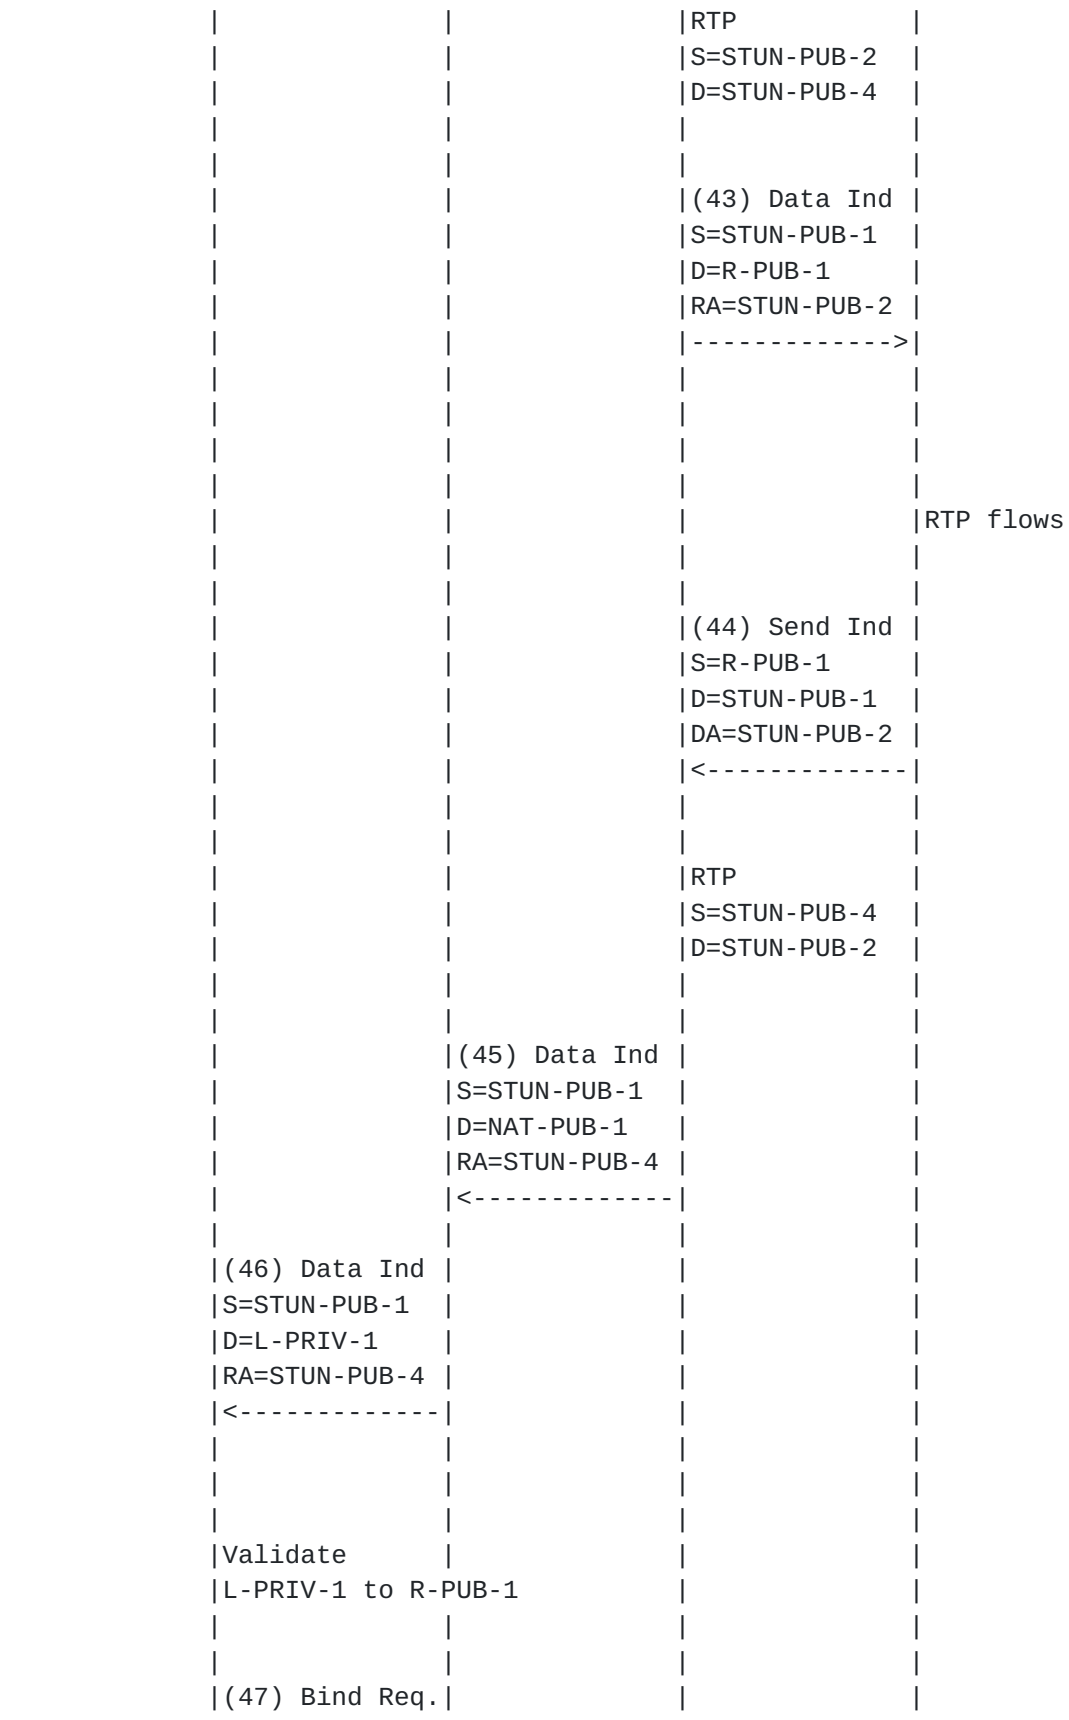
RTP
S=STUN-PUB-2
D=STUN-PUB-4
(43) Data Ind
S=STUN-PUB-1
D=R-PUB-1
RA=STUN-PUB-2
RTP flows
(44) Send Ind
S=R-PUB-1
D=STUN-PUB-1
DA=STUN-PUB-2
RTP
S=STUN-PUB-4
D=STUN-PUB-2
(45) Data Ind
S=STUN-PUB-1
D=NAT-PUB-1
RA=STUN-PUB-4
(46) Data Ind
S=STUN-PUB-1
D=L-PRIV-1
RA=STUN-PUB-4
Validate
L-PRIV-1 to R-PUB-1
(47) Bind Req.

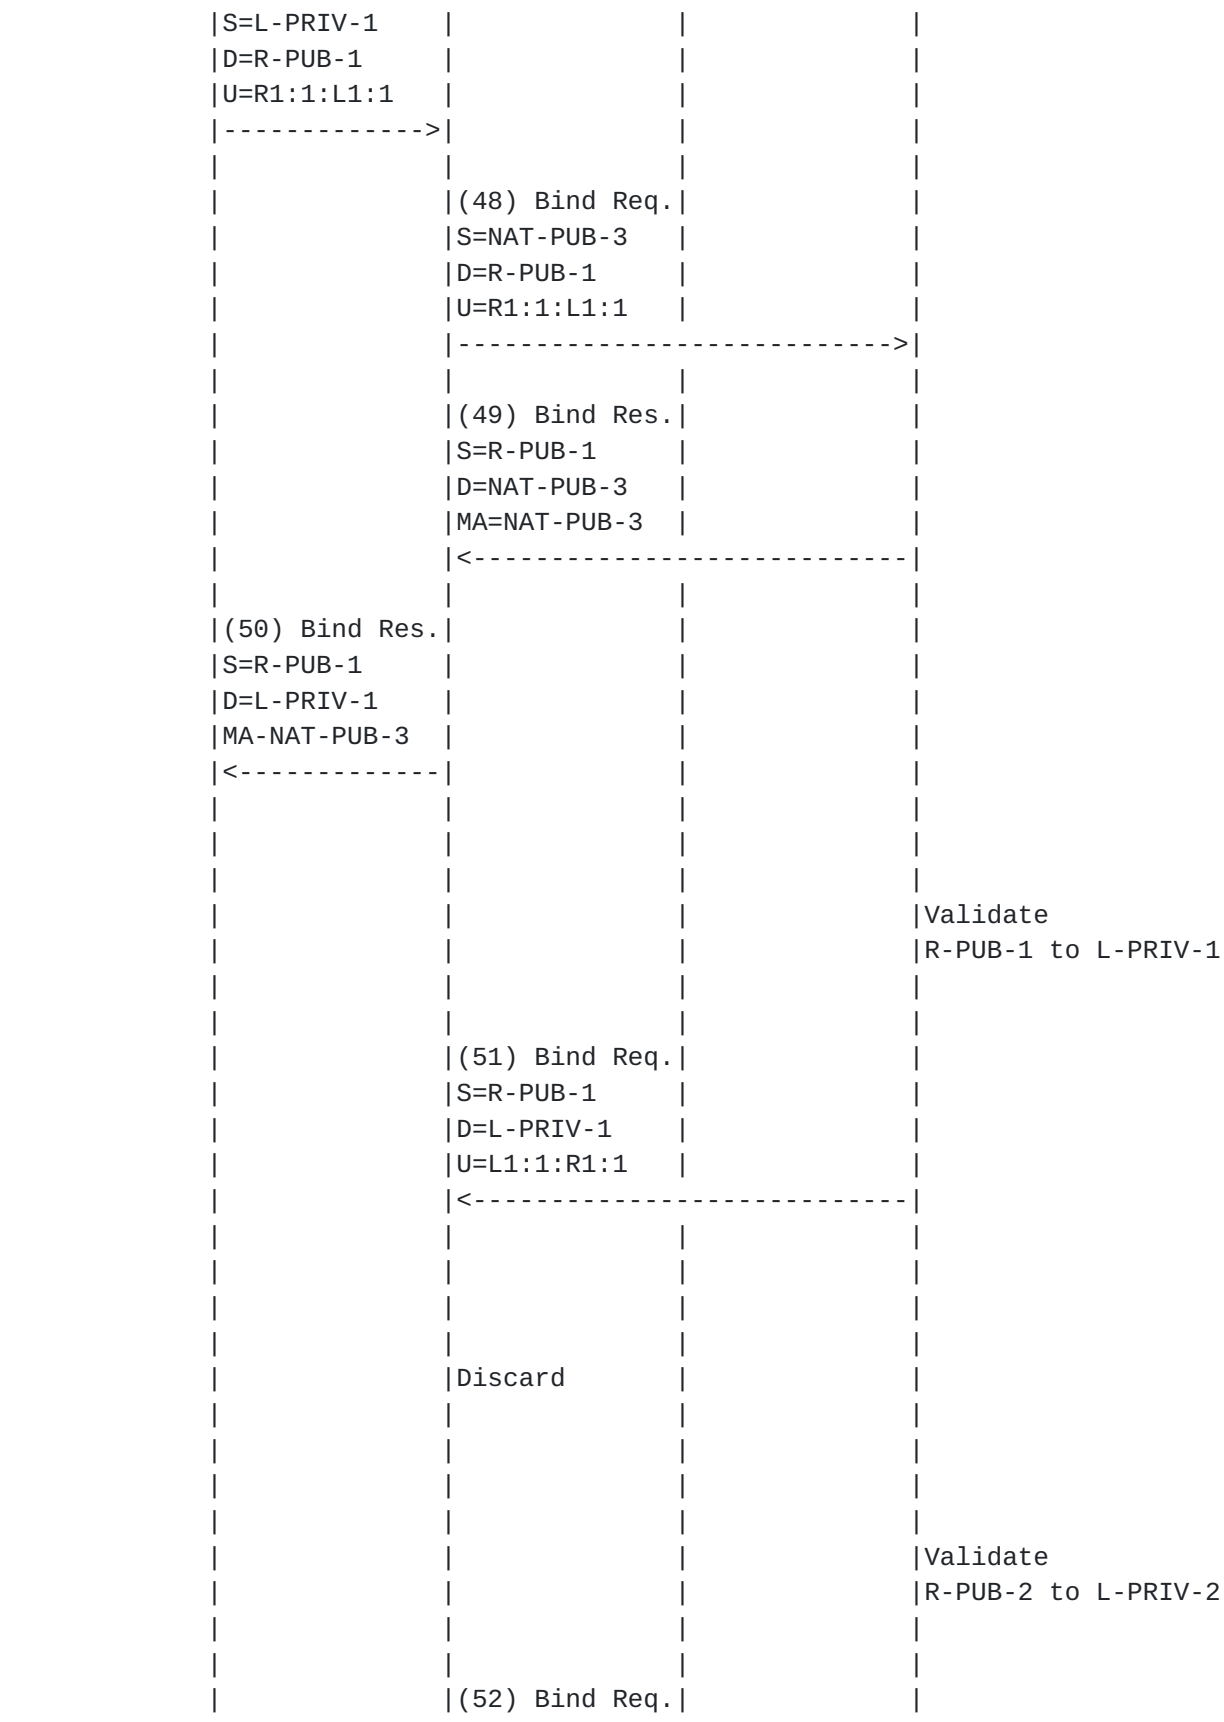
S=L-PRIV-1
D=R-PUB-1
U=R1:1:L1:1
(48) Bind Req.
S=NAT-PUB-3
D=R-PUB-1
U=R1:1:L1:1
(49) Bind Res.
S=R-PUB-1
D=NAT-PUB-3
MA=NAT-PUB-3
(50) Bind Res.
S=R-PUB-1
D=L-PRIV-1
MA-NAT-PUB-3
Validate
R-PUB-1 to L-PRIV-1
(51) Bind Req.
S=R-PUB-1
D=L-PRIV-1
U=L1:1:R1:1
Discard
Validate
R-PUB-2 to L-PRIV-2
(52) Bind Req.

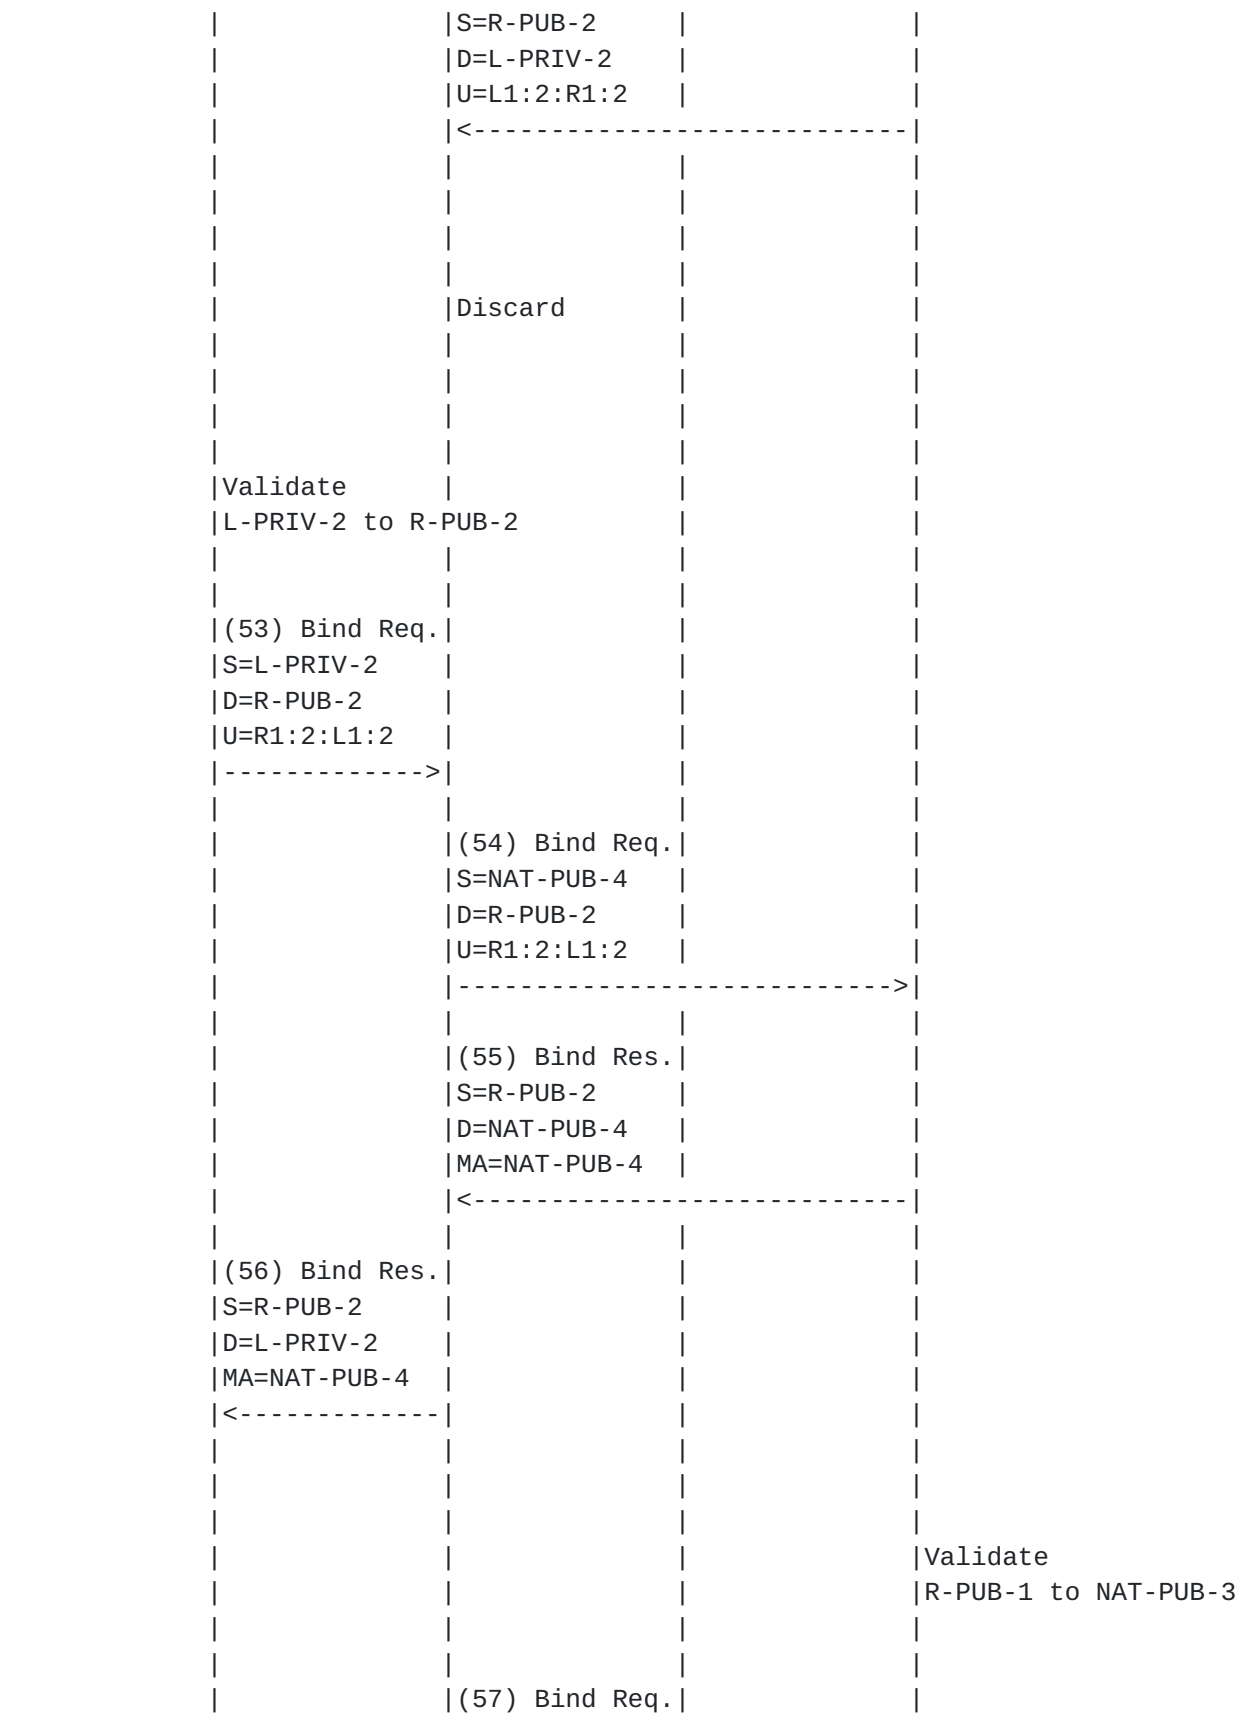
S=R-PUB-2
D=L-PRIV-2
U=L1:2:R1:2
Discard
Validate
L-PRIV-2 to R-PUB-2
(53) Bind Req.
S=L-PRIV-2
D=R-PUB-2
U=R1:2:L1:2
(54) Bind Req.
S=NAT-PUB-4
D=R-PUB-2
U=R1:2:L1:2
(55) Bind Res.
S=R-PUB-2
D=NAT-PUB-4
MA=NAT-PUB-4
(56) Bind Res.
S=R-PUB-2
D=L-PRIV-2
MA=NAT-PUB-4
Validate
R-PUB-1 to NAT-PUB-3
(57) Bind Req.

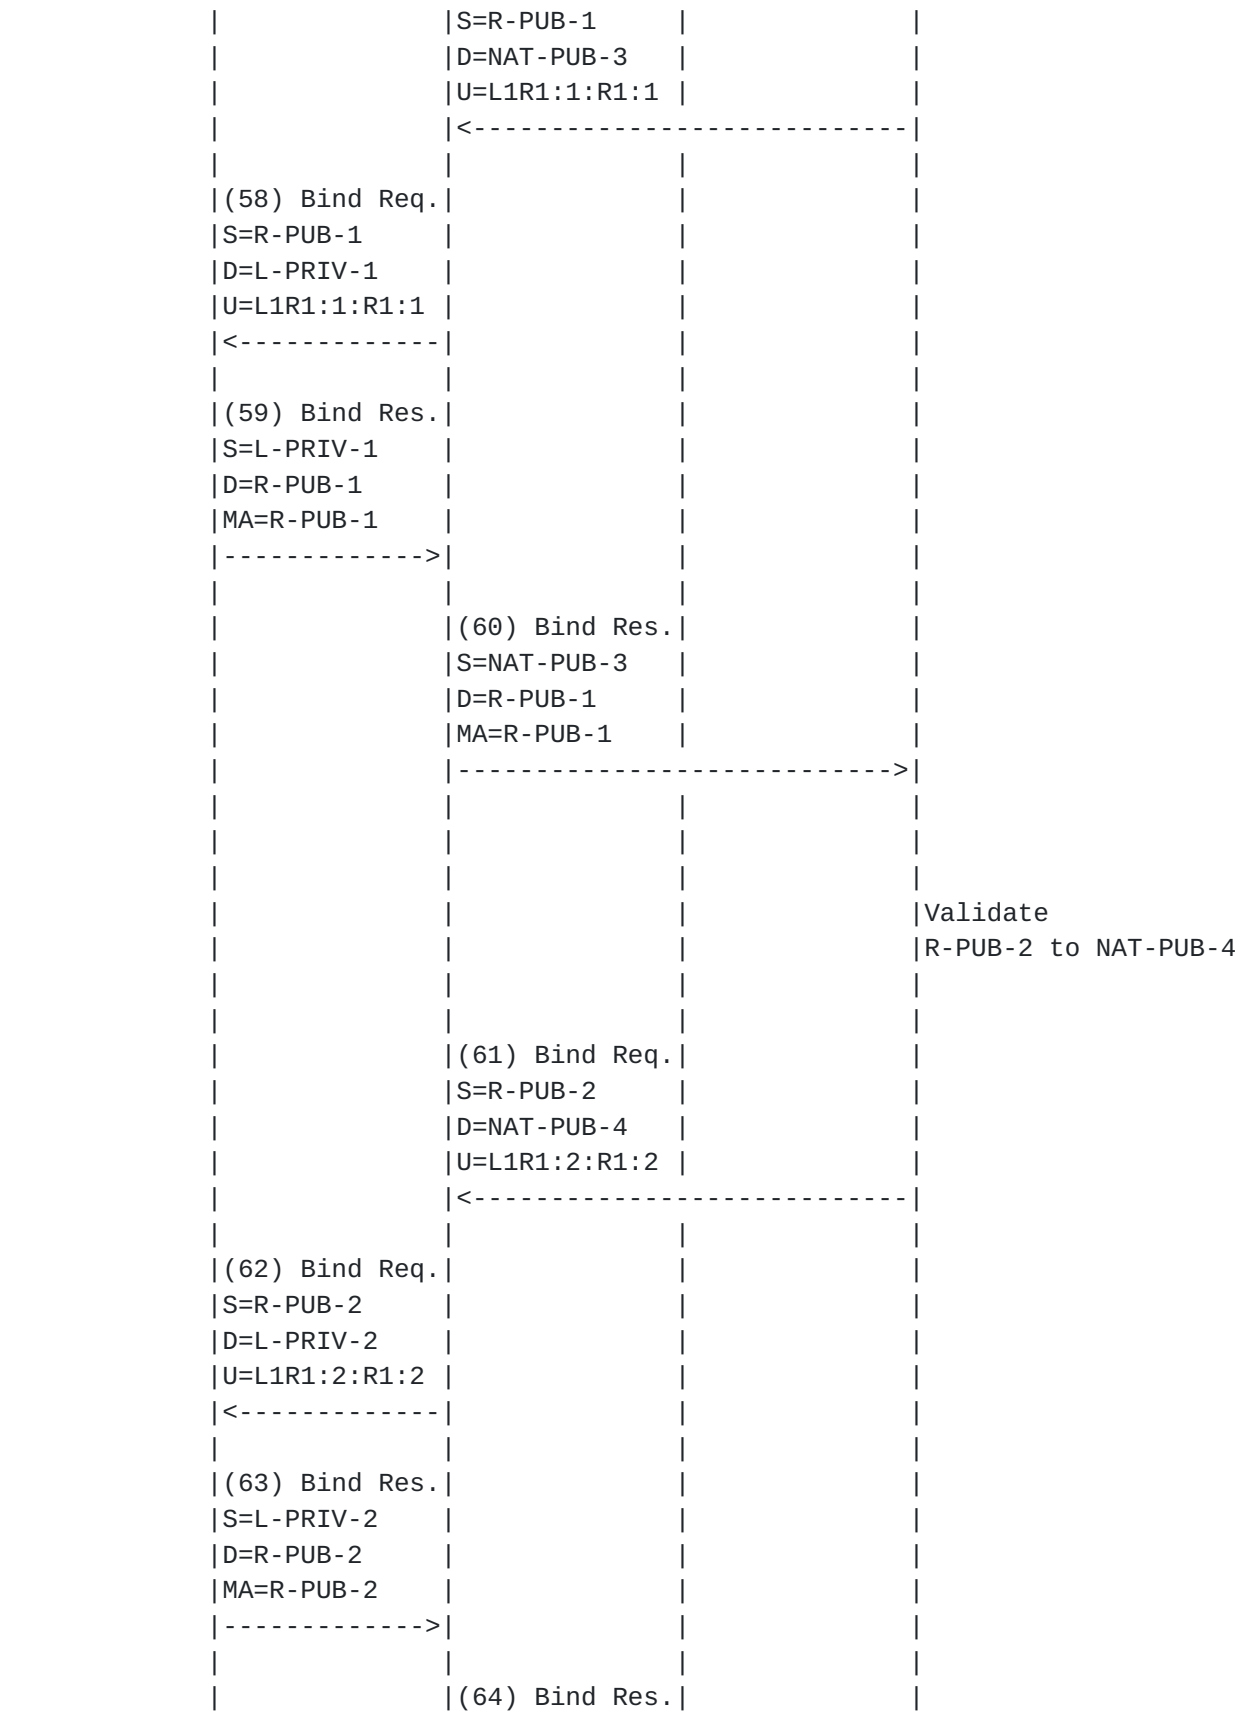
S=R-PUB-1
D=NAT-PUB-3
U=L1R1:1:R1:1
(58) Bind Req.
S=R-PUB-1
D=L-PRIV-1
U=L1R1:1:R1:1
(59) Bind Res.
S=L-PRIV-1
D=R-PUB-1
MA=R-PUB-1
(60) Bind Res.
S=NAT-PUB-3
D=R-PUB-1
MA=R-PUB-1
Validate
R-PUB-2 to NAT-PUB-4
(61) Bind Req.
S=R-PUB-2
D=NAT-PUB-4
U=L1R1:2:R1:2
(62) Bind Req.
S=R-PUB-2
D=L-PRIV-2
U=L1R1:2:R1:2
(63) Bind Res.
S=L-PRIV-2
D=R-PUB-2
MA=R-PUB-2
(64) Bind Res.

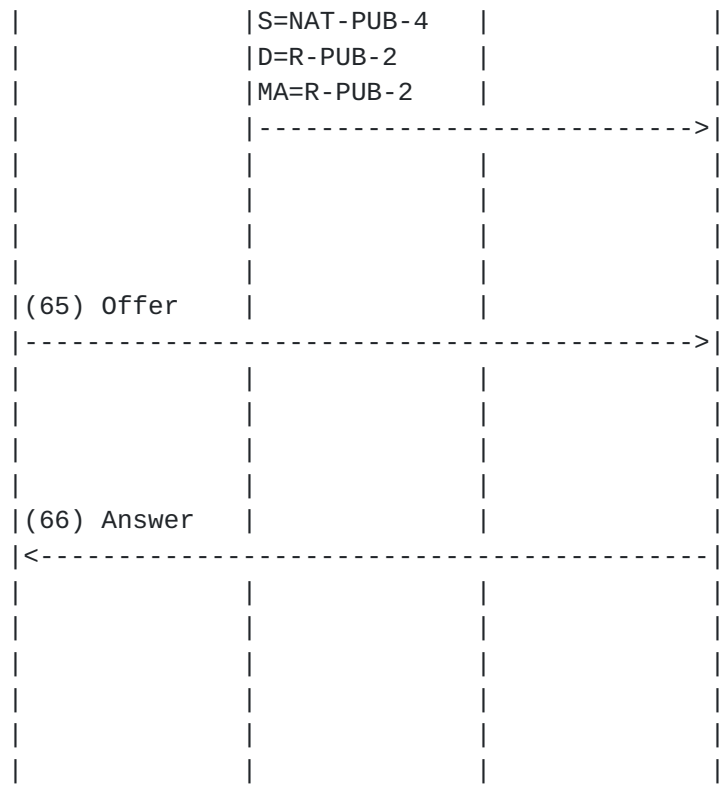

Figure 11

First, agent L obtains both server reflexive and relayed transport addresses for its RTP packets, using a STUN Allocate request, which will provide it with both types of addresses (messages 1-4). Recall that the NAT has the address and port dependent mapping property. Here, it creates a binding of NAT-PUB-1 for this UDP request, and this becomes the server reflexive transport address for RTP. The relayed transport address is STUN-PUB-2, allocated by the STUN server. Agent L repeats this process for RTCP (messages 5-8) Ta seconds later, and obtains NAT-PUB-2 as its server reflexive transport address for RTCP and STUN-PUB-3 for its relayed transport address.

With its three candidates, agent L prioritizes them, choosing the local candidate as highest priority, followed by the server reflexive candidate, followed by the relayed candidate. It chooses its relayed candidate as the active candidate, and encodes it into the m/c-line. The resulting offer (message 17) looks like:

```
v=0
o=jdoe 2890844526 2890842807 IN IP4 $L-PRIV-1.IP
s=
c=IN IP4 $STUN-PUB-2.IP
t=0 0
a=ice-pwd:$LPASS
m=audio $STUN-PUB-2.PORT RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=rtcp:$STUN-PUB-3.PORT
a=candidate $L1 1 UDP 1.0 $L-PRIV-1.IP $L-PRIV-1.PORT
a=candidate $L1 2 UDP 1.0 $L-PRIV-2.IP $L-PRIV-2.PORT
a=candidate $L2 1 UDP 0.7 $NAT-PUB-1.IP $NAT-PUB-1.PORT
a=candidate $L2 2 UDP 0.7 $NAT-PUB-2.IP $NAT-PUB-2.PORT
a=candidate $L3 1 UDP 0.3 $STUN-PUB-2.IP $STUN-PUB-2.PORT
a=candidate $L3 2 UDP 0.3 $STUN-PUB-3.IP $STUN-PUB-3.PORT
```

This offer is received at agent R. Agent R will gather its server reflexive and relayed transport addresses for RTP from an Allocate request (messages 10-11). Since the server reflexive transport address matches its local transport address, no separate candidate is used for it. The agent then gathers its server reflexive and relayed transport addresses for RTCP (messages 12-13). It prioritizes the local candidate with higher priority than the relayed candidate, and selects the relayed candidate as the active candidate. Its resulting answer looks like:

```
v=0
o=bob 2808844564 2808844564 IN IP4 $R-PUB-1.IP
s=
c=IN IP4 $STUN-PUB-4.IP
t=0 0
a=ice-pwd:$RPASS
m=audio $STUN-PUB-4.PORT RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=rtcp:$STUN-PUB-5.PORT
a=candidate $R1 1 UDP 1.0 $R-PUB-1.IP $R-PUB-1.PORT
a=candidate $R1 2 UDP 1.0 $R-PUB-2.IP $R-PUB-2.PORT
a=candidate $R2 1 UDP 0.3 $STUN-PUB-4.IP $STUN-PUB-4.PORT
a=candidate $R2 2 UDP 0.3 $STUN-PUB-5.IP $STUN-PUB-5.PORT
```

Next, agents L and R form candidate pairs and the transport address check ordered list. This list will start with the two components in the currently active candidate pair - relayed candidates. Agent R begins its checks (message 15). It will check connectivity between the active candidate pair, starting with the first component, which is STUN-PUB-4 for agent R and STUN-PUB-2 for agent L. The state machine for that transport address pair moves to the Testing state.

Since this is a relayed transport address for agent R, it utilizes the STUN Send Indication to deliver the Binding Request. The DESTINATION-ADDRESS is STUN-PUB-2.

The STUN server will extract the content of the Send indication, which is a STUN Binding Request, and deliver it to the destination, STUN-PUB-4. This request will be sent from the relayed address allocated to R, which is STUN-PUB-4. As both interfaces are on the STUN server, this message is sent to itself (and thus the lack of a message number in the sequence diagram above). Note that the USERNAME in the Binding Request is L3:1:R2:1, which represents the transport address pair ID. This message gets discarded by the STUN server since, as of yet, there are no permissions established for the STUN-PUB-2 allocation. However, it did have the side effect of establishing a permission on the STUN-PUB-4 binding, allowing incoming packets from STUN-PUB-2.

Once L gets the offer, it will attempt to validate the first transport address pair in the transport address pair check ordered list, which will be the active candidate. The state machine for this transport address pair moves into the Testing state. Like agent R did, it will use the STUN Send Indication to send a STUN Binding Request from its relayed transport address, STUN-PUB-2, to STUN-PUB-4 (message 16). This packet traverses the NAT (message 17) and arrives at the STUN server. The STUN server will unwrap the contents of the packet and send them from STUN-PUB-2 to STUN-PUB-4. It will also, as a consequence, add a permission for STUN-PUB-4. The contents of the packet are a STUN Binding Request with USERNAME R2:1:L3:1 (note how this is the flip of the USERNAME in the Binding Request sent by agent R). This is also a packet from the STUN server to itself. However, now, the packet is not discarded, as a permission had been installed as a consequence of the "suicide packet" from agent R (a suicide packet is a packet that has no hope of traversing a far end NAT, but serves the purpose of enabling a permission in a near end NAT so that a packet from the peer can be returned). Thus, the STUN server will relay the received STUN request towards agent R (message 18). This is delivered as a STUN Data Indication. Notice how the REMOTE-ADDRESS is STUN-PUB-2; this is important as it will be used to construct the STUN Binding Response.

Agent R will receive the Data Indication, and unwrap its contents to find the Binding Request. The state machine for this transport address pair is currently in the Testing state. It therefore moves into the Send-Valid state, and it generates a Binding Response. However, the XOR-MAPPED-ADDRESS in the Binding Response is constructed using the source IP address and port that were seen by the STUN server when the Binding Request arrived at STUN-PUB-4, which is the looped message between messages 17 and 18. This source

address is STUN-PUB-2, which is the value of the REMOTE-ADDRESS attribute in message 18. Thus, the STUN Binding Response will contain STUN-PUB-2 in the XOR-MAPPED-ADDRESS, and is to be sent to STUN-PUB-2. To send the response, agent R takes the STUN Binding Response and encapsulates it in a STUN Send indication, setting the DESTINATION-ADDRESS to STUN-PUB-2. This is shown in message 19.

The STUN server will receive this Send Indication, and unwrap its contents to find the STUN Binding Response. It sends it to the value of the DESTINATION-ADDRESS attribute, and sends it from the relayed address allocated to R, which is STUN-PUB-4. This, once again, results in a looped message to itself, and it arrives at STUN-PUB-2. Now, however, there is a permission installed for STUN-PUB-4. The STUN server will therefore forward the packet to agent L. To do so, it constructs a STUN Data Indication containing the contents of the packet. It sets the REMOTE-ADDRESS to the source transport address of the request it received (STUN-PUB-4), and forwards it to agent L (message 20). This traverses the NAT (message 21) and arrives at agent L. As a consequence of the receipt of a Binding Response, the state machine for this transport address pair moves to the Recv-Valid state. The agent also examines the XOR-MAPPED-ADDRESS of the STUN response. It indicates STUN-PUB-2. This is the same as the native transport address of this transport address pair, and thus doesn't represent a new transport address that might have been learned.

Because of the receipt of message 18, the transport address pair moved from Testing to Send-Valid, causing R to attempt a retransmission of its STUN Binding Request that was lost (the contents of message 15 that were discarded by the STUN server due to lack of permission). This time, however, a permission has been installed and the retransmission will work. So, it sends the Binding Request again (message 22, identical to message 15). This is looped by the STUN server to itself again, but this time there is a permission in place when it arrives at STUN-PUB-2. As such, the request is forwarded towards agent L this time, in a STUN Data Indication (message 23). This traverses the NAT (message 24) and arrives at agent L. Agent L extracts the contents of the request, which are a STUN Binding Request. This causes the state machine to move from Recv-Valid to Valid. It generates a STUN Binding Response, and sets the XOR-MAPPED-ADDRESS based on the value of the REMOTE-ADDRESS in message 24 (STUN-PUB-4). This Binding Response is sent to STUN-PUB-4, which is accomplished through a STUN Send Indication (message 25). This Send Indication traverses the NAT (message 26) and is received by the STUN server. Its contents are decapsulated, and sent to STUN-PUB-4, which is again a loop on the same host. This packet is then sent towards agent R in a Data Indication (message 27). The contents of the DATA Indication are extracted, and the agent sees a successful Binding Response. It therefore moves the

state machine from the Send-Valid state to the Valid state. At this point, the transport address pair is in the Valid state for both agents.

Approximately Ta seconds after agent R sent message 15, agent R will start checks for the next transport address pair in its transport address pair check ordered list. This is the second component of the same candidate pair, used for RTCP. This sequence, messages 28 through 40, are identical to the ones for RTP, but differ only in the specific transport addresses.

Once that validation happens, the second transport address pair has been validated. The candidate pair moves into the valid state, and both candidates are considered valid. The active candidate has now been validated, and media can begin to flow. It will do so through the STUN server; indeed, it is relayed "twice" through the STUN server. Even though there is a single STUN server, it is logically acting as two separate STUN servers. Indeed, had L and R used two separate STUN servers, media would be relayed through both STUN servers in a trapezoid configuration.

The actual media flows are shown as well. It is important to note that, since the ICE checks have not yet concluded on the candidate that will ultimately be used, no STUN Set Active Destinations have been sent. As a consequence, media that is sent through the STUN servers has to be sent using STUN Send indications. This introduces some overhead, but is a transient condition. In message 41, agent L sends an RTP packet to agent R using a Send indication. It is sent to STUN-PUB-4. This traverses the NAT (message 42), and arrives at the STUN server. It is decapsulated, looped to itself, and arrives at STUN-PUB-4. From there, it is encapsulated in a Data Indication and sent to agent R (message 43). In the reverse direction, agent R will send an RTP packet using a STUN Send indication (message 42), and send it to STUN-PUB-2. This is received by the STUN server, decapsulated, and sent to STUN-PUB-2 from STUN-PUB-4. This is again a loop within the same host, arriving at STUN-PUB-4. The contents of the packet are sent to agent L through a STUN Data Indication (message 45), which traverses the NAT (message 46) to arrive at agent L. Since this call flow is already long enough, RTCP packet transmission is not shown.

Approximately Ta seconds after it sends message 29, agent L goes to the next transport address pair in its transport address pair check ordered list that is in the Waiting state. This will be the RTP candidate for the top priority candidate pair, which is L-PRIV-1 on agent L and R-PUB-1 on agent R. This is a local candidate for each agent. To perform the check, agent L sends a STUN Binding Request from L-PRIV-1 to R-PUB-1 (message 47). Note the USERNAME of

R1:1:L1:1, which identifies this transport address pair. This traverses the NAT (message 48). Since the NAT has the address and port dependent mapping property, and this is a new destination IP address, the NAT allocates a new transport address on its public side, NAT-PUB-3, and places this in the source IP address and port. This packet arrives at agent R. Agent R finds a matching transport address pair in the Waiting state. The state machine transitions to the Send-Valid state. It sends the Binding response, with a XOR-MAPPED-ADDRESS indicating NAT-PUB-3 (message 49), which traverses the NAT and arrives at agent L (message 50). Agent R, in addition to sending the response, will also send a Binding Request. It is important to remember that this Binding Request is sent to the remote address in the transport address pair (L-PRIV-1), and NOT to the source IP address and port of the Binding Request (NAT-PUB-3); that will happen later. This attempt is shown in message 51. However, since the L-PRIV-1 is private, the packet is discarded in the network.

Now, as a consequence of receiving message 48, agent R will have constructed a peer-derived candidate. The candidate ID for this candidate is L1R1, and it initially contains a single transport address pair, NAT-PUB-3 and R-PUB-1. However, the candidate isn't yet usable until the other component gets added. Similarly, agent L will have constructed the same peer-derived candidate, with the same candidate ID and the same transport address pair.

Some Ta seconds after sending message 28, agent R will move to the next transport address pair in the transport address pair check ordered list whose state is Waiting. This is the RTCP component of the highest priority candidate pair. It will attempt a connectivity check, from R-PUB-2 to L-PRIV-2 (message 52). Since L-PRIV-1 is private, this message is discarded.

Some Ta seconds after sending message 47, agent L will move to the next transport address pair in the transport address pair check ordered list whose state is Waiting. This is the RTCP component of the highest priority candidate pair. It will attempt a connectivity check, from L-PRIV-2 to R-PUB-2 (message 53), which operates nearly identically to messages 47-50, with the exception of the specific addresses. Here, the NAT will create a new binding for the RTCP, NAT-PUB-4, and this transport address is new for both participants. On receipt of this Binding Request at agent R (message 54), agent R constructs the candidate ID for the peer-derived candidate, L1R1, and finds it already exists. As such, this new transport address is added, and the peer-derived candidate becomes complete and usable. Agent L does the same thing on receipt of message 56. This candidate will have the same priority as its generating candidate L1 (1.0), and is paired up with R1 (also at priority 1.0). Since L1R1 has the same

priority as L1 itself, the ordering algorithm in Section 7.5 will use the reverse lexicographic order of the candidate ID iself to determine order. L1R1 is larger than L1, so that the peer-derived candidate will come before its generating candidate. As a consequence, the peer-derived candidate pair will have a higher priority than its generating candidate, and appear just before it in the candidate pair priority ordered list.

As a consequence, after agent R sends message 55 and completes the peer-derived candidate, it will move the two transport addresses in the peer derived candidate into the Send-Valid state, and send a Binding Request for each in rapid succession (agent L will have moved both into the Recv-Valid state upon receipt of message 56). The first of these connectivity checks are for the RTP component, from R-PUB-1 to NAT-PUB-3 (message 57). Note the USERNAME in the STUN Binding Request, L1R1:1:R1:1, which identifies the peer-derived transport address pair. This will succesfully traverse the NAT and be delivered to agent L (message 58). The receipt of this request moves the state machine for this transport address pair from Recv-Valid to Valid, and a Binding Response is sent (message 59). This passes through the NAT and arrives at agent R (message 60). This causes its state machine to enter the Valid state as well. The reflexive transport address, R-PUB-1, is not new to agent R and thus does not result in the creation of a new peer-derived candidate.

Messages 61 through 64 show the same basic flow for RTCP. Upon receipt of message 64, both transport address pairs are Valid at both agents, causing the peer derived candidate to become valid. Timer Tws is set at agent L, and fires without any higher priority candidate pairs becoming validated. At agent R, media can now be sent on this candidate pair from answerer (agent R) to offerer (agent L). Agent L sends an updated offer to promote the peer-derived candidate to active. This offer (message 65) looks like:

```
v=0
o=jdoe 2890844526 2890842808 IN IP4 $L-PRIV-1.IP
s=
c=IN IP4 $NAT-PUB-3.IP
t=0 0
a=ice-pwd:$LPASS
m=audio $NAT-PUB-3.PORT RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=rtcp:$NAT-PUB-4.PORT
a=remote-candidate:R1
a=candidate $L1 1 UDP 1.0 $L-PRIV-1.IP $L-PRIV-1.PORT
a=candidate $L1 2 UDP 1.0 $L-PRIV-2.IP $L-PRIV-2.PORT
```

There are several important things to note in this offer. Firstly, note how the m/c-line now contains NAT-PUB-3 and NAT-PUB-4, the peer derived transport addresses it learned through the ICE processing. Secondly, note how there remains a candidate encoded into the a=candidate attributes. This is candidate L1, NOT candidate L1R1. Recall that the peer-derived candidates are never encoded into the SDP. Rather, their generating candidate is encoded. This will cause keepalives to take place for the generating candidate if valid (though its not) and any of its derived candidates, which is what we want. Finally, notice the inclusion of the a=remote-candidate attribute. Since agent L doesn't know whether agent R received messages 60 or 64, it doesnt know whether the state of the candidate is Send-Valid or Valid at agent R. So, it has to tell agent R that, in case its Send-Valid, to please use it anyway.

The answer generated by agent R looks like:

```
v=0
o=bob 2808844564 2808844565 IN IP4 $R-PUB-1.IP
s=
c=IN IP4 $R-PUB-1.IP
t=0 0
a=ice-pwd:$RPASS
m=audio $R-PUB-1.PORT RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=rtcp:$R-PUB-2.PORT
a=candidate $R1 1 UDP 1.0 $R-PUB-1.IP $R-PUB-1.PORT
a=candidate $R1 2 UDP 1.0 $R-PUB-2.IP $R-PUB-2.PORT
```

With this, media can now flow directly between endpoints. The removal of the relayed candidates from the offer/answer exchange will cause the STUN relay allocations to be removed.

12. Grammar

This specification defines three new SDP attributes - the "candidate", "remote-candidate" and "ice-pwd" attributes.

The candidate attribute is a media-level attribute only. It contains a transport address for a candidate that can be used for connectivity checks. There may be multiple candidate attributes in a media block.

The syntax of this attribute is defined using Augmented BNF as defined in RFC 4234 [9]:

```
candidate-attribute   = "candidate" ":" candidate-id SP component-id SP
                        transport SP
                        qvalue SP   ;qvalue from RFC 3261
                        addr SP     ;addr from RFC 3266
                        port        ;port from RFC 2327
                        *(SP extension-att-name SP
                             extension-att-value)

transport             = "UDP" / transport-extension
transport-extension   = token
candidate-id          = 1*base64-char

base64-char           = ALPHANUM / DIGIT / "+" / "/"
                              ;ALPHANUM from RFC 3261
component-id          = 1*DIGIT
extension-att-name    = byte-string    ;from RFC 2327
extension-att-value   = byte-string
```

The candidate-id is used to group together the transport addresses for a particular candidate. It MUST be constructed with at least 24 bits of randomness. It MUST have the same value for all transport addresses within the same candidate. It MUST have a different value for transport addresses within different candidates for the same media stream. The candidate-id uses a syntax that is defined to be equal to the base64 alphabet [3], which allows the candidate-id to be generated by performing a base64 encoding of a randomly generated value (note, however, that this does not mean that the candidate-id or password is base64 decoded when use in STUN messages). In addition, if content is base64 encoded to generate the candidate-id, it MUST NOT be padded with '='. Section 2.2 of RFC 3548 indicates that some base64 usages do not require padding, and it requests that such usages call out that fact. ICE is one such usage. This is because the data is never decoded. The component-id is a positive integer, which identifies the specific component of the candidate. It MUST start at 1 and MUST increment by 1 for each component of a particular candidate.

The addr production is taken from [10], allowing for IPv4 addresses, IPv6 addresses and FQDNs. The port production is taken from RFC 2327 [5]. The token production is taken from RFC 3261 [2]. The transport production indicates the transport protocol for the candidate. This specification only defines UDP. However, extensibility is provided to allow for future transport protocols to be used with ICE, such as TCP or the Datagram Congestion Control Protocol (DCCP) [34].

The a=candidate attribute can itself be extended. The grammar allows for new name/value pairs to be added at the end of the attribute. An

implementation MUST ignore any name/value pairs it doesn't understand.

The syntax of the "remote-candidate" attribute is defined using Augmented BNF as defined in RFC 4234 [9]:

```
remote-candidate-att = "remote-candidate" ":" candidate-id
```

This attribute MUST be present in an offer when the candidate in the m/c-line is part of a candidate pair that is in the valid or partially valid state.

The syntax of the "ice-pwd" attribute is defined as:

```
ice-pwd-att           = "ice-pwd" ":" password
password              = 1*base64-char
```

The "ice-pwd" attribute can appear at either the session-level or media-level. When present in both, the value in the media-level takes precedence. Thus, the value at the session level is effectively a default that applies to all media streams, unless overriden by a media-level value. It MUST have at least 128 bits of randomness. Like the candidate-ID, its syntax is taken from the base64 alphabet, allowing the password to be generted from a base64 encoding of a 128 bit value. In addition, if content is base64 encoded to generate the candidate-id, it MUST NOT be padded with '='.

13. Security Considerations

There are several types of attacks possible in an ICE system. This section considers these attacks and their countermeasures.

13.1 Attacks on Connectivity Checks

An attacker might attempt to disrupt the STUN-based connectivity checks. Ultimately, all of these attacks fool an agent into thinking something incorrect about the results of the connectivity checks. The possible false conclusions an attacker can try and cause are:

False Invalid: An attacker can fool a pair of agents into thinking a candidate pair is invalid, when it isn't. This can be used to cause an agent to prefer a different candidate (such as one injected by the attacker), or to disrupt a call by forcing all candidates to fail.

False Valid: An attacker can fool a pair of agents into thinking a candidate pair is valid, when it isn't. This can cause an agent to proceed with a session, but then not be able to receive any media.

False Peer-Derived Candidate: An attacker can cause an agent to discover a new peer-derived candidate, when it shouldn't have. This can be used to redirect media streams to a DoS target or to the attacker, for eavesdropping or other purposes.

False Valid on False Candidate: An attacker has already convinced an agent that there is a candidate with an address that doesn't actually route to that agent (for example, by injecting a false peer-derived candidate or false STUN-derived candidate). It must then launch an attack that forces the agents to believe that this candidate is valid.

Of the various techniques for creating faked STUN messages described in [13], many are not applicable for the connectivity checks. Compromises of STUN servers are not much of a concern, since the STUN servers are embedded in endpoints and distributed throughout the network. Thus, compromising the STUN server is equivalent to comprimising the endpoint, and if that happens, far more problematic attacks are possible than those against ICE. Similarly, DNS attacks are irrelevant since STUN servers are not discovered via DNS, they are signaled via SIP. Injection of fake responses and relaying modified requests all can be handled in ICE with the countermeasures discussed below.

To force the false invalid result, the attacker has to wait for the connectivity check for one of the agents to be sent. When it is, the attacker needs to inject a fake response with an unrecoverable error response, such as a 600. This attack only needs to be launched against one of the agents in order to invalidate the candidate pair. However, since the candidate is, in fact, valid, the original request may reach the peer agent, and result in a success response. The attacker needs to force this packet or its response to be dropped, through a DoS attack, layer 2 network disruption, or other technique. If it doesn't do this, the success response will also reach the originator, alerting it to a possible attack. This will cause the agent to abandon the candidate, which is the desired result in any case. Fortunately, this attack is mitigated completely through the STUN message integrity mechanism. The attacker needs to inject a fake response, and in order for this response to be processed, the attacker needs the password. If the offer/answer signaling is secured, the attacker will not have the password.

Forcing the fake valid result works in a similar way. The agent

needs to wait for the Binding Request from each agent, and inject a fake success response. The attacker won't need to worry about disrupting the actual response since, if the candidate is not valid, it presumably wouldn't be received anyway. However, like the fake invalid attack, this attack is mitigated completely through the STUN message integrity and offer/answer security techniques.

Forcing the false peer-derived candidate result can be done either with fake requests or responses, or with replays. We consider the fake requests and responses case first. It requires the attacker to send a Binding Request to one agent with a source IP address and port for the false transport address. In addition, the attacker must wait for a Binding Request from the other agent, and generate a fake response with a XOR-MAPPED-ADDRESS attribute. This attack is best launched against a candidate pair that is likely to be invalid, so the attacker doesnt need to contend with the actual responses to the real connectivity checks. Like the other attacks described here, this attack is mitigated by the STUN message integrity mechanisms and secure offer/answer exchanges.

Forcing the false peer-derived candidate result with packet replays is different. The attacker waits until one of the agents sends a Binding Request for one of the transport address pairs. It then intercepts this request, and replays it towards the other agent with a faked source IP address. It must also prevent the original request from reaching the remote agent, either by launching a DoS attack to cause the packet to be dropped, or forcing it to be dropped using layer 2 mechanisms. The replayed packet is received at the other agent, and accepted, since the integrity check passes (the integrity check cannot and does not cover the source IP address and port). It is then responded to. This response will contain a XOR-MAPPED-ADDRESS with the false transport address. It is passed to the this false address. The attacker must then intercept it and relay it towards the originator.

The other agent will then initiate a connectivity check towards that transport address. This validation needs to succeed. This requires the attacker to force a false valid on a false candidate. Injecting of fake requests or responses to achieve this goal is prevented using the integrity mechanisms of STUN and the offer/answer exchange. Thus, this attack can only be launched through replays. To do that, the attacker must intercept the Binding Request towards this false transport address, and replay it towards the other agent. Then, it must intercept the response and replay that back as well.

This attack is very hard to launch unless the attacker themself is identified by the fake transport address. This is because it requires the attacker to intercept and replay packets sent by two

different hosts. If both agents are on different networks (for example, across the public Internet), this attack can be hard to coordinate, since it needs to occur against two different endpoints on different parts of the network at the same time.

If the attacker themself is identified by the fake transport address, the attack is easier to coordinate. However, if SRTP is used [24], the attacker will not be able to play the media packets, they will only be able to discard them, effectively disabling the media stream for the call. However, this attack requires the agent to disrupt packets in order to block the connectivity check from reaching the target. In that case, if the goal is to disrupt the media stream, its much easier to just disrupt it with the same mechanism, rather than attack ICE.

13.2 Attacks on Address Gathering

ICE endpoints make use of STUN for gathering addresses from a STUN server in the network. This is corresponds to the binding acquisition use case discussed in Section 10.1 of [13]. As a consequence, the attacks against STUN itself that are described in Section 12 [13] can still be used against the STUN address gathering operations that occur in ICE.

However, the additional mechanisms provided by ICE actually counteract such attacks, making binding acquisition with STUN more secure when combined with ICE than without ICE.

Consider an attacker which is able to provide an agent with a faked XOR-MAPPED-ADDRESS in a STUN Binding Request that is used for address gathering. This is the primary attack primitive described in Section 12 of [13]. This address will be used as a STUN derived candidate in the ICE exchange. For this candidate to actually be used for media, the attacker must also attack the connectivity checks, and in particular, force a false valid on a false candidate. This attack is very hard to launch if the false address identifies a third party, and is prevented by SRTP if it identifies the attacker themself.

If the attacker elects not to attack the connectivity checks, the worst it can do is prevent the STUN-derived address from being used. However, if the peer agent has at least one address that is reachable by the agent under attack, the STUN connectivity checks themselves will provide a STUN-derived address that can be used for the exchange of media. Peer derived candidates are preferred over the candidate they are generated from for this reason. As such, an attack solely on the STUN address gathering will normally have no impact on a call at all.

13.3 Attacks on the Offer/Answer Exchanges

An attacker that can modify or disrupt the offer/answer exchanges themselves can readily launch a variety of attacks with ICE. They could direct media to a target of a DoS attack, they could insert themselves into the media stream, and so on. These are similar to the general security considerations for offer/answer exchanges, and the security considerations in RFC 3264 [4] apply. These require techniques for message integrity and encryption for offers and answers, which are satisfied by the SIPS mechanism [2] when SIP is used. As such, the usage of SIPS with ICE is RECOMMENDED.

13.4 Insider Attacks

In addition to attacks where the attacker is a third party trying to insert fake offers, answers or stun messages, there are several attacks possible with ICE when the attacker is an authenticated and valid participant in the ICE exchange.

13.4.1 The Voice Hammer Attack

The voice hammer attack is an amplification attack, of the variety discussed in Section 3 of [32]. In this attack, the attacker initiates sessions to other agents, and includes the IP address and port of a DoS target in the m/c-line of their SDP. This causes substantial amplification; a single offer/answer exchange can create a continuing flood of media packets, possibly at high rates (consider video sources). This attack is not speific to ICE, but ICE can help provide remediation.

Specifically, if ICE is used, the agent receiving the malicious SDP will first peform connectivity checks to the target of media before sending it there. If this target is a third party host, the checks will not succeed, and media is never sent.

Unfortunately, ICE doesn't help if its not used, in which case an attacker could simply send the offer without the ICE parameters. However, in environments where the set of clients are known, and limited to ones that support ICE, the server can reject any offers or answers that don't indicate ICE support.

13.4.2 STUN Amplification Attack

The STUN amplification attack is similar to the voice hammer. However, instead of voice packets being directed to the target, STUN connectivity checks are directed to the target. This attack is accomplished by having the offerer send an offer with a large number of candidates, say 50. The answerer receives the offer, and starts

its checks, which are directed at the target, and consequently, never generate a response. The answerer will start a new connectivity check every 50ms, and each check is a STUN transaction consisting of 9 retransmits of a message 64 bytes in length. This produces a fairly substantial 92 kbps, just in STUN requests.

It is impossible to eliminate the amplification, but the volume can be reduced through a variety of heuristics. For example, agents can limit the number of candidates they'll accept in an offer or answer, they can increase the value of Ta, or exponentially increase Ta as time goes on. All of these ultimately trade off the time for the ICE exchanges to complete, with the amount of traffic that gets sent.

14. IANA Considerations

This specification defines three new SDP attribute per the procedures of Appendix B of RFC 2327. The required information for the registrations are included here.

14.1 candidate Attribute

Contact Name: Jonathan Rosenberg, jdrosen@jdrosen.net.

Attribute Name: candidate

Long Form: candidate

Type of Attribute: media level

Charset Considerations: The attribute is not subject to the charset attribute.

Purpose: This attribute is used with Interactive Connectivity Establishment (ICE), and provides one of many possible candidate addresses for communication. These addresses are validated with an end-to-end connectivity check using Simple Traversal of UDP with NAT (STUN).

Appropriate Values: See Section 12 of RFC XXXX [Note to RFC-ed: please replace XXXX with the RFC number of this specification].

14.2 remote-candidate Attribute

Contact Name: Jonathan Rosenberg, jdrosen@jdrosen.net.

Attribute Name: remote-candidate

Long Form: remote-candidate

Type of Attribute: media level

Charset Considerations: The attribute is not subject to the charset attribute.

Purpose: This attribute is used with Interactive Connectivity Establishment (ICE), and provides the identity of the remote candidate that the offerer wishes the answerer to use in its answer.

Appropriate Values: See Section 12 of RFC XXXX [Note to RFC-ed: please replace XXXX with the RFC number of this specification].

14.3 ice-pwd Attribute

Contact Name: Jonathan Rosenberg, jdrosen@jdrosen.net.

Attribute Name: ice-pwd

Long Form: ice-pwd

Type of Attribute: session level

Charset Considerations: The attribute is not subject to the charset attribute.

Purpose: This attribute is used with Interactive Connectivity Establishment (ICE), and provides the password used to protect STUN connectivity checks.

Appropriate Values: See Section 12 of RFC XXXX [Note to RFC-ed: please replace XXXX with the RFC number of this specification].

15. IAB Considerations

The IAB has studied the problem of "Unilateral Self Address Fixing", which is the general process by which a agent attempts to determine its address in another realm on the other side of a NAT through a collaborative protocol reflection mechanism [22]. ICE is an example of a protocol that performs this type of function. Interestingly, the process for ICE is not unilateral, but bilateral, and the difference has a signficant impact on the issues raised by IAB. The

IAB has mandated that any protocols developed for this purpose document a specific set of considerations. This section meets those requirements.

15.1 Problem Definition

From RFC 3424 any UNSAF proposal must provide:

> Precise definition of a specific, limited-scope problem that is to be solved with the UNSAF proposal. A short term fix should not be generalized to solve other problems; this is why "short term fixes usually aren't".

The specific problems being solved by ICE are:

> Provide a means for two peers to determine the set of transport addresses which can be used for communication.
>
> Provide a means for resolving many of the limitations of other UNSAF mechanisms by wrapping them in an additional layer of processing (the ICE methodology).
>
> Provide a means for a agent to determine an address that is reachable by another peer with which it wishes to communicate.

15.2 Exit Strategy

From RFC 3424, any UNSAF proposal must provide:

> Description of an exit strategy/transition plan. The better short term fixes are the ones that will naturally see less and less use as the appropriate technology is deployed.

ICE itself doesn't easily get phased out. However, it is useful even in a globally connected Internet, to serve as a means for detecting whether a router failure has temporarily disrupted connectivity, for example. ICE also helps prevent certain security attacks which have nothing to do with NAT. However, what ICE does is help phase out other UNSAF mechanisms. ICE effectively selects amongst those mechanisms, prioritizing ones that are better, and deprioritizing ones that are worse. Local IPv6 addresses can be preferred. As NATs begin to dissipate as IPv6 is introduced, derived transport addresses from other UNSAF mechanisms simply never get used, because higher priority connectivity exists. Therefore, the servers get used less and less, and can eventually be remove when their usage goes to zero.

Indeed, ICE can assist in the transition from IPv4 to IPv6. It can

be used to determine whether to use IPv6 or IPv4 when two dual-stack hosts communicate with SIP (IPv6 gets used). It can also allow a network with both 6to4 and native v6 connectivity to determine which address to use when communicating with a peer.

15.3 Brittleness Introduced by ICE

From RFC3424, any UNSAF proposal must provide:

> Discussion of specific issues that may render systems more "brittle". For example, approaches that involve using data at multiple network layers create more dependencies, increase debugging challenges, and make it harder to transition.

ICE actually removes brittleness from existing UNSAF mechanisms. In particular, traditional STUN (as described in [16]) has several points of brittleness. One of them is the discovery process which requires a agent to try and classify the type of NAT it is behind. This process is error-prone. With ICE, that discovery process is simply not used. Rather than unilaterally assessing the validity of the address, its validity is dynamically determined by measuring connectivity to a peer. The process of determining connectivity is very robust.

Another point of brittleness in STUN and any other unilateral mechanism is its absolute reliance on an additional server. ICE makes use of a server for allocating unilateral addresses, but allows agents to directly connect if possible. Therefore, in some cases, the failure of a STUN server would still allow for a call to progress when ICE is used.

Another point of brittleness in traditional STUN is that it assumes that the STUN server is on the public Internet. Interestingly, with ICE, that is not necessary. There can be a multitude of STUN servers in a variety of address realms. ICE will discover the one that has provided a usable address.

The most troubling point of brittleness in traditional STUN is that it doesn't work in all network topologies. In cases where there is a shared NAT between each agent and the STUN server, traditional STUN may not work. With ICE, that restriction can be lifted.

Traditional STUN also introduces some security considerations. Fortunately, those security considerations are also mitigated by ICE.

Consequently, ICE serves to repair the brittleness introduced in other UNSAF mechanisms, and does not introduce any additional brittleness into the system.

15.4 Requirements for a Long Term Solution

From RFC 3424, any UNSAF proposal must provide:

> Identify requirements for longer term, sound technical solutions -- contribute to the process of finding the right longer term solution.

Our conclusions from STUN remain unchanged. However, we feel ICE actually helps because we believe it can be part of the long term solution.

15.5 Issues with Existing NAPT Boxes

From RFC 3424, any UNSAF proposal must provide:

> Discussion of the impact of the noted practical issues with existing, deployed NA[P]Ts and experience reports.

A number of NAT boxes are now being deployed into the market which try and provide "generic" ALG functionality. These generic ALGs hunt for IP addresses, either in text or binary form within a packet, and rewrite them if they match a binding. This interferes with traditional STUN. However, the update to STUN [13] uses an encoding which hides these binary addresses from generic ALGs. Since [13] is required for all ICE implementations, this NAPT problem does not impact ICE.

Existing NAPT boxes have non-deterministic and typically short expiration times for UDP-based bindings. This requires implementations to send periodic keepalives to maintain those bindings. ICE uses a default of 15s, which is a very conservative estimate. Eventually, over time, as NAT boxes become compliant to behave [37], this minimum keepalive will become deterministic and well-known, and the ICE timers can be adjusted. Having a way to discover the minimum keepalive interval would be far better still.

16. Acknowledgements

The authors would like to thank Flemming Andreasen, Rohan Mahy, Dean Willis, Dan Wing, Douglas Otis, Tim Moore, and Francois Audet for their comments and input. A special thanks goes to Magnus Westerlund for doing several detailed reviews on the various revisions of this specification. His input led to many substantive improvements in this document.

17. References

17.1 Normative References

[1] Huitema, C., "Real Time Control Protocol (RTCP) attribute in Session Description Protocol (SDP)", RFC 3605, October 2003.

[2] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

[3] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 3548, July 2003.

[4] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.

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

[6] Casner, S., "Session Description Protocol (SDP) Bandwidth Modifiers for RTP Control Protocol (RTCP) Bandwidth", RFC 3556, July 2003.

[7] Camarillo, G., Marshall, W., and J. Rosenberg, "Integration of Resource Management and Session Initiation Protocol (SIP)", RFC 3312, October 2002.

[8] Camarillo, G. and P. Kyzivat, "Update to the Session Initiation Protocol (SIP) Preconditions Framework", RFC 4032, March 2005.

[9] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", RFC 4234, October 2005.

[10] Olson, S., Camarillo, G., and A. Roach, "Support for IPv6 in Session Description Protocol (SDP)", RFC 3266, June 2002.

[11] Rosenberg, J. and H. Schulzrinne, "Reliability of Provisional Responses in Session Initiation Protocol (SIP)", RFC 3262, June 2002.

[12] Yon, D., "Connection-Oriented Media Transport in the Session Description Protocol (SDP)", draft-ietf-mmusic-sdp-comedia-10 (work in progress), November 2004.

[13] Rosenberg, J., "Simple Traversal of UDP Through Network Address Translators (NAT) (STUN)", draft-ietf-behave-rfc3489bis-03 (work in progress), March 2006.

[14] Rosenberg, J., Mahy, R., and C. Huitema, "Obtaining Relay

Addresses from Simple Traversal of UDP Through NAT (STUN)", Internet Draft draft-ietf-behave-turn-00.txt, February 2006.

17.2 Informative References

[15] Schulzrinne, H., Rao, A., and R. Lanphier, "Real Time Streaming Protocol (RTSP)", RFC 2326, April 1998.

[16] Rosenberg, J., Weinberger, J., Huitema, C., and R. Mahy, "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)", RFC 3489, March 2003.

[17] Senie, D., "Network Address Translator (NAT)-Friendly Application Design Guidelines", RFC 3235, January 2002.

[18] Rosenberg, J. and H. Schulzrinne, "An RTP Payload Format for Generic Forward Error Correction", RFC 2733, December 1999.

[19] Srisuresh, P., Kuthan, J., Rosenberg, J., Molitor, A., and A. Rayhan, "Middlebox communication architecture and framework", RFC 3303, August 2002.

[20] Borella, M., Lo, J., Grabelsky, D., and G. Montenegro, "Realm Specific IP: Framework", RFC 3102, October 2001.

[21] Borella, M., Grabelsky, D., Lo, J., and K. Taniguchi, "Realm Specific IP: Protocol Specification", RFC 3103, October 2001.

[22] Daigle, L. and IAB, "IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation", RFC 3424, November 2002.

[23] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", RFC 3550, July 2003.

[24] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.

[25] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.

[26] Zopf, R., "Real-time Transport Protocol (RTP) Payload for Comfort Noise (CN)", RFC 3389, September 2002.

[27] Rosenberg, J., "The Session Initiation Protocol (SIP) UPDATE Method", RFC 3311, October 2002.

[28] Bonica, R., Kompella, K., and D. Meyer, "Tracing Requirements for Generic Tunnels", RFC 3609, September 2003.

[29] Camarillo, G. and H. Schulzrinne, "Early Media and Ringing Tone Generation in the Session Initiation Protocol (SIP)", RFC 3960, December 2004.

[30] Andreasen, F., "Connectivity Preconditions for Session Description Protocol Media Streams", draft-ietf-mmusic-connectivity-precon-01 (work in progress), October 2005.

[31] Andreasen, F., "A No-Op Payload Format for RTP", draft-ietf-avt-rtp-no-op-00 (work in progress), May 2005.

[32] Rescorla, E. and M. Handley, "Internet Denial of Service Considerations", draft-iab-dos-03 (work in progress), September 2005.

[33] Huitema, C., "Teredo: Tunneling IPv6 over UDP through NATs", draft-huitema-v6ops-teredo-05 (work in progress), April 2005.

[34] Kohler, E., "Datagram Congestion Control Protocol (DCCP)", draft-ietf-dccp-spec-13 (work in progress), December 2005.

[35] Lazzaro, J., "Framing RTP and RTCP Packets over Connection-Oriented Transport", draft-ietf-avt-rtp-framing-contrans-06 (work in progress), September 2005.

[36] Hellstrom, G., "RTP Payload for Text Conversation", draft-ietf-avt-rfc2793bis-09 (work in progress), August 2004.

[37] Audet, F. and C. Jennings, "NAT Behavioral Requirements for Unicast UDP", Internet Draft draft-ietf-behave-nat-udp-00.txt, February 2006.

Author's Address

Jonathan Rosenberg
Cisco Systems
600 Lanidex Plaza
Parsippany, NJ 07054
US

Phone: +1 973 952-5000
Email: jdrosen@cisco.com
URI: http://www.jdrosen.net

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.