Document link:
https://tools.ietf.org/html/draft-ietf-ice-rfc5245bis-12

Summary

Almost there. As with any document, it could be improved.
Most important points:

- The protocol has been designed and revised to be usable with non-media data, but the introduction and abstract do not reflect this. Expunging the media bias from the body of the document is probably not worth it, but the intro and abstract should mention it.
- Security considerations should mention the problem that ICE reveals addresses that might otherwise remain hidden, and that this is a privacy concern.
- The document has removed all the SDP specific parts (good), but the requirements it places on the negotiation mechanism aren't collectively documented anywhere. A section describing this would help comprehension for people developing signalling protocols for use with ICE.
- The definition of “component” talks about a component having one address. I believe that in current usage, it should be defined to have an address pair. (non-symmetric RTP is dead).

The rest of my suggested changes are nits, I think.

Abstract

Is it really correct to cite this protocol as being only for multimedia? Suggestion: replace “multimedia” with “communication”. This is in line with bullet 1 of section 21 - “generalized”.

Section 1 - Introduction

Similar to the above: In order to reduce the scope of changes, suggest adding to the end of the introduction:
“This specification may be used both for multimedia sessions and other types of communication.”

Section 3 - Terminology

Terminology: “Component”: “A component is a piece of a media stream requiring a single transport address” - I think this should be “a single pair of transport addresses”, or perhaps even use the defined term “candidate pair”.

Suggested reformulation: “A component is a piece of a media stream that has to be assigned a single candidate pair, which can be used by no other component, before being used for transmitting part of a media stream”.

“User”, as used in section 4.1.1 “prior to alerting the user”, is not defined.

Section 4 - ICE candidate gathering and exchange

4.1 - part of this section title is missing. It may have been intended to be “Full Implementation”, since 4.2 is “Lite”.
4.1.1 - nit: “Every candidate is a transport address” -> “has a transport address”.

4.1.1 - should one mention that candidates don’t live forever? (NAT timeout, interfaces going up or down, network mobility)

4.1.1.1 - should “Host candidates corresponding to IPv6 link-local addresses MUST NOT be gathered” be its own bullet, possibly placed last with a condition that these are only gathered if there is no global IPv6 address on the same interface? Or is this strictly in order to prevent lcoation tracking?

4.1.1.1 - temporary and permanent IPv6 addresses are not mentioned. Draft-ietf-rtcweb-transports section 3.3 recommends discarding permanent addresses, as per RFC 6724. Should this advice be added to this section?

4.1.1.2 - paragraph 3 ends “Allocate requests SHOULD be authenticated using a” - something’s missing. RFC 5245 had “long-term credential obtained by the client through some other means.”

Paragraph 4 - first occurence of “Ta”. Please expand. (Suggest adding a paragraph: The gathering process is controlled using a timer called “Ta”).

4.1.1.4 - host candidates die too. Suggest adding a paragraph: “Host candidates don’t time out, but the addresses may change or disappear for a number of reasons. The agent SHOULD monitor the interfaces it uses, invalidate candidates whose base has gone away, and acquire new candidates as appropriate when new interfaces appear.”

4.1.2 grammar nit - “since both agents will not be coordinated in their checks” -> “since the agents will not be coordinated in their checks”.

4.1.2.1 obsolete reference to RFC 3484 for “precedence value for IP addresses”. Note that the default precedence table in 6724 section 2.1 is quite different from the one in 3484, so switching the reference will give different candidate orderings. (Also, 6724 is referenced by section 4.2, so this is likely an oversight.)

4.1.2.2 should the paragraph mention that VPN candidates and local LAN candidates can have different local preferences too? I think their order is application (or even user) dependent, but the user's preference is likely to be quite strong (when using VPN for corp-LAN connectivity: don't use VPN when LAN connection is available - when using VPN for location hiding: don't use host at all when VPN is available - but details are out of scope for this specification)

4.3 nit: Since the section says that encoding is out of scope, the title should be changed to "Exchanging the Candidate Information".

4.3 "Related Address" is not defined previously. A pointer to section B.3 would be appropriate. (It seems equally undefined there what is supposed to be sent. This seems doubtful for interoperability. Should it be punted to a higher layer, as in "Use or non-use of the Related Address is defined by the ICE usage"?)

4.4 Grammar nit - "may alter the ICE candidate information that breaks ICE" -> "may alter the ICE candidate information in ways that break ICE".

Section 5. ICE Candidate Processing

5.1.1 determining role - the "Both lite" bullet is the only one that mentions that both agents will believe they are controlled or controlling. I think this is true for both case 1 and 3.

I believe this paragraph belongs at a higher level, and should be reformulated as:
"The signaling protocol enabling the candidate exchange MUST allow the agents to determine which of the agents is the Initiating Agent, and whether or not the peer agent is Full or Lite". This can also replace the second NOTE, and avoids reference to "glare" (that's the signalling protocol's problem).
(I don't see why one shouldn't allow the 6.3.1.1 procedure to apply in this case too, in which case the signalling protocol requirement can be a SHOULD, with fallback to 6.3.1.1, which removes the need to call out 3pcc as a separate case. But that may be a too large change at this stage.)

5.1.2.1 grammar nit: "is neither Completed yet nor Failed yet" - remove the first "yet".
5.1.2.2 spell: para 1 - educed -> reduced

5.1.2.6 state procedure, bullet 5: "in the first list (ordered by the check list set)" doesn't make sense. "In the first list according to the usage-defined check list set order" would make sense. The order is ICE usage controlled per bullet 2 of the same list.

5.1.4.1 This paragraph would be clearer if it added the words "The queue is initially empty, and will have candidates entered as part of the procedures in section 6."

5.2 Grammar nit “On determining …. means …” is ungrammatical. Suggested rephrasing: “If the lite implementation is the controlling agent (which will only happen if the peer agent is also a lite implementation), it selects ….. and then updates ….”

Section 6. Performing connectivity checks

Nit: Intro paragraph: doubled “to to”.

6.2.3. Diffserv treatment - doesn’t address the case of media wishing to utilize multiple DSCP markings in its media stream (draft-ietf-tsvwg-rtcweb-qos, table 1). Suggest adding “If multiple DSCP markings are used on the media packets, the agent SHOULD choose one of them for use with the connectivity check.”

6.2.5.2. Failure - nit: “as a result a connectivity check” - missing “of”

Sections 6.3.1.3 “Learning Peer Reflexive Candidates” and section 6.3.1.4 “Triggered checks” don’t form separate procedures, the way the other subsections of 6.3.1 do. I suggest they be merged.
6.2.5.2.2 describes processing ICMP errors for connectivity checks. ICMP errors are unauthenticated and easily spoofable; using them for shutting down communication is a well known vulnerability.

Timeout would seem to be the normal outcome of a connectivity check; it seems strange to hide this most common case under “Unrecoverable STUN response”. Suggest raising it to its own section.

6.2.5.3 Success doesn’t say “The following sections (6.2.5.3.1 to 6.2.5.3.4) are performed in order”. I think it should, since other sections have subsections (6.2.5 is an example) where only one is chosen based on some criteria.

Section 6.2.5.3.2 introduces the “valid list”, which is one list for the whole check list set. Section 6.3.2 introduces the “valid list” as if it was a new concept. There should be only one, and its definition probably needs to be lifted higher - it’s also referenced from section 7.

Section 7 - Concluding ICE processing

Section 7.1.1 starts using CHECK LIST and VALID LIST in uppercase, but inconsistently. Inconsistency is bad; consider lowercasing them all.

Nit: nomiation -> nomination (para 4). Need a comma before “and for selecting a pair”.

“Ice2” ice option needs a forward reference, since this is the first mention, and definition comes later.

Section 10 - Keepalives

The text is unclear about whether “An agent MAY use another value for Tr” is allowed to override the “MUST NOT use a Tr value smaller than 15 seconds”. I think it doesn’t (has to be more than 15), but I’m not 100% that’s the intended meaning.

This spec says that all endpoints MUST send keepalives, but doesn’t say which ones. Draft-ietf-rtcweb-stun-consent-freshness specifies a 30-second usage of STUN bind requests, for instance, but it’s not clear whether it uses STUN keepalives or not. Should this document say explicitly that the ICE Usage specifies which keepalives are used?

Section 12 - Receiving media

“.... if there is a change in source transport address, but the media packets come from the same peer agent, this SHOULD NOT be treated as an SSRC collision”. Why not MUST NOT?

Section 14 setting Ta and RTO

14.2 Ta - “the agent MUST indicate the proposed value to its peer during the ICE exchange”. The term “ICE exchange” is never defined; I think it refers to the exchange of candidates via the signalling protocol. Please define the term.

Section 16 - Security Considerations:

There is no privacy consideration - in particular, the privacy impact of exposing host-specific addresses together with VPN addresses is not mentioned or addressed.
Suggestion: Add a section saying “The process of probing for candidates reveals the source addresses of the client and its peer to any on-network listening attacker, and the process of exchanging candidates reveals the addresses to any attacker that is able to see the negotiation. Some addresses, such as the server reflexive addresses gathered through the local interface of VPN users, may be sensitive information. If these potential attacks can’t be mitigated, the implementation may want to institute controls for which addresses are revealed to the negotiation and/or probing process. Such controls need to be specified as part of the ICE usage.”