Multihoming L3 Shim Approach

draft-ietf-multi6-l3shim-00.txt
Erik Nordmark
erik.nordmark@sun.com
Marcelo Bagnulo
marcelo@it.uc3m.es

Agenda

- Changes since previous version
  - draft-nordmark-multi6dt-shim-00.txt
- Future changes

Changes (1)

- Added assumption that something else handles the interaction between ingress filtering and source address selection

- Clarified things with respect to using ULAs in general, and added separate text about centrally assigned ULAs

- Added more text about MTU dropping implications and ICMP too big re-mapping

- Added text specifying how the shim handles the flow label field, and the impact on flow setup protocols

Changes (2)

- Added text about the need for the sender to handle ICMP errors

- Added text that there might be other protocols than flow setup protocols and ICMP errors that might be impacted by the shim

- Added text about IP multicast in a new section

- Added clarification in section 8.4 about AAAA records being for a service and not a host, needing some care

Changes (3)

- Added a clarification in section 8.4 that learning the different locators during initial communication from the DNS potentially has different trust issues than learning them from the peer.

- Clarified the two models of flow label usage for demultiplexing

- In section 5 clarified that state maintenance is not per ULP connection

- In section 5 clarified merging option

Changes (4)

- Clarified in section 9.1 why it isn't sufficient to avoid using the same locators for different ULIDs for the same peer host

- Clarified in section 9.1.1/9.1.2 that there is multi6 state at the receiver to tell how/whether to rewrite the source address field.

- Clarified the aspect of section 9.2.1 which talks about not being able to use a new locator until the peer has been told of the new locator.

Changes (5)

- Added text about the implications of renumbering and reassignment.

- Clarified section on flow labels to

  - first talk about the simple case of using <source locator, destination locator, flow label> and its complexities, and

  - then about the potential to just use the flow label by itself to identify the context.

TODO (1)

- Using "address" vs. "locator" and "ULID" more consistently and carefully

- Q: whether the interaction between source locator selection and ingress filtering implies a stronger assumption

  - A: It might be too early to make that strong assumption

- Make it clear that the probability of prefix reuse causing address reuse it very small, so it might be overkill to stop using a ULID when it becomes invalid

TODO (2)

- Point out that MTU change can occur from locator pair switching, and not only from adding an extension header

- More clarifications on what is ULA specific vs. just related to the reverse DNS tree

Next steps

- Are there other issues/comments?

- Issue 01 with the TODO changes above