

Suggested clarifications and updates to the Babel TLV processing

Denis Ovsienko (Custodian Data Centres)

denis@custodiandc.com

IETF 97 Seoul, November 18 2016

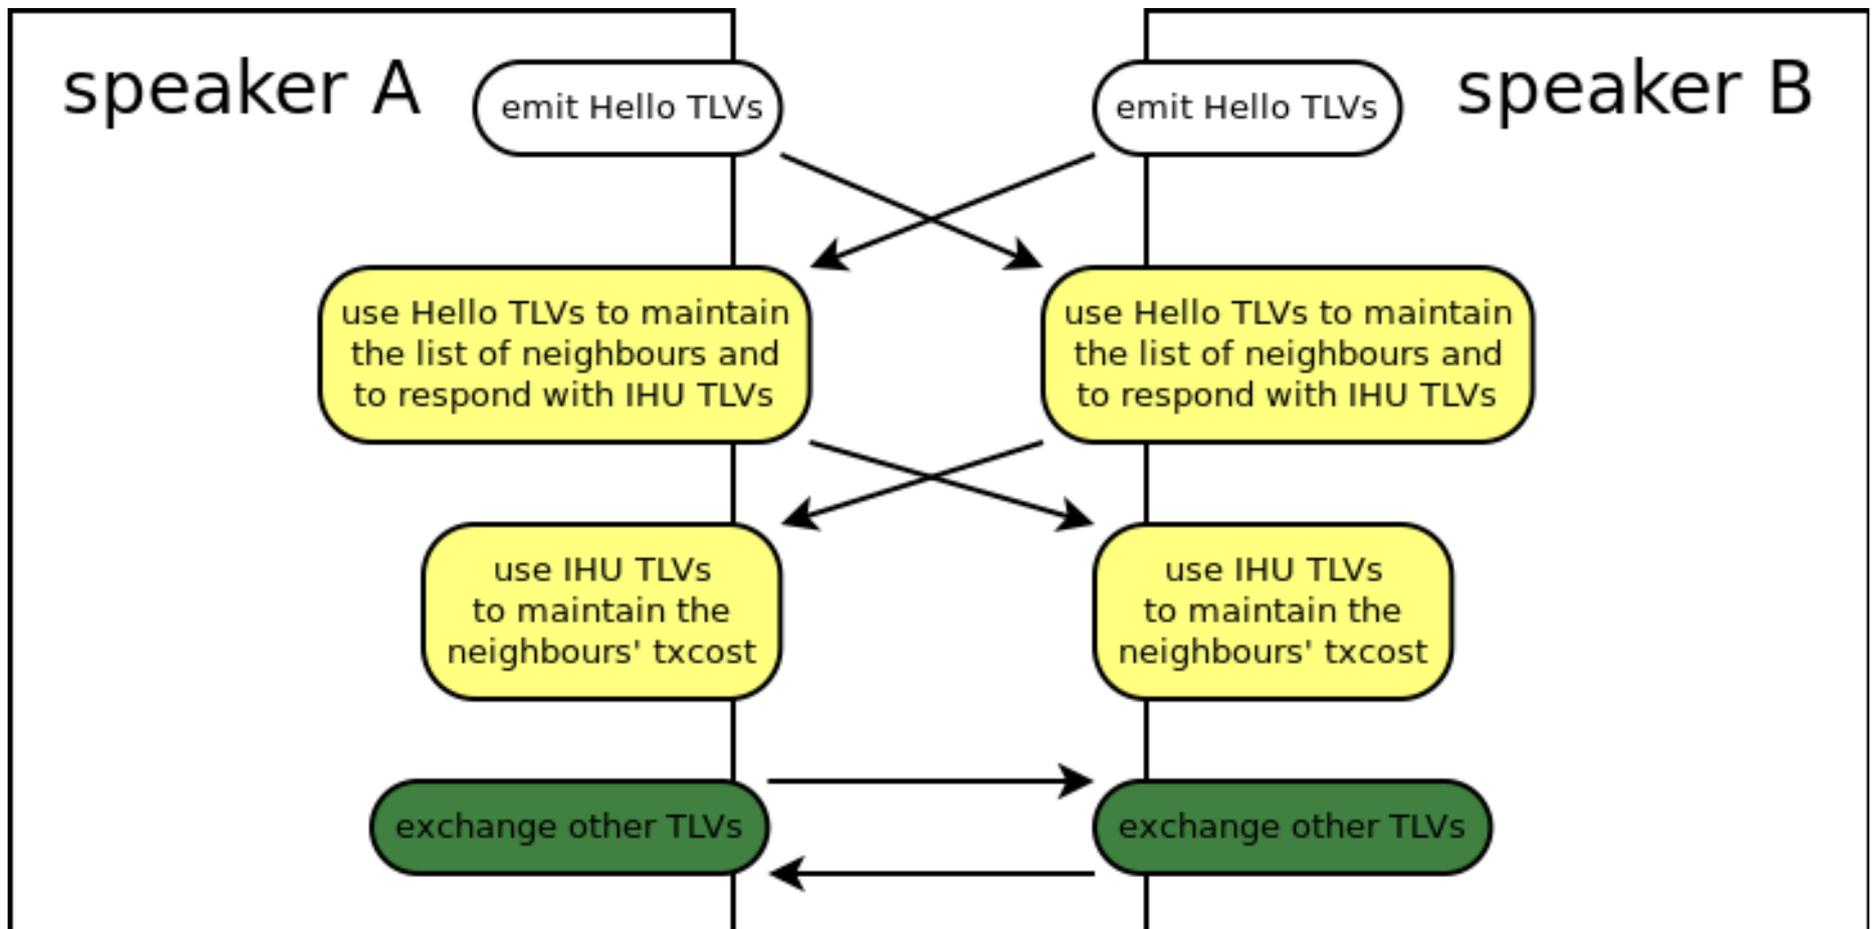
I. Suggested clarifications to the existing normative text in RFC 6126-bis regarding bidirectional neighbour reachability in Babel

The normative text concerned

- Section 3.4.1: Reverse Reachability Detection
- Section 3.4.2: Bidirectional Reachability Detection
- Appendix A.1: Maintaining Hello History

The next slide tries to represent the above combined.

Looks reasonable at a glance

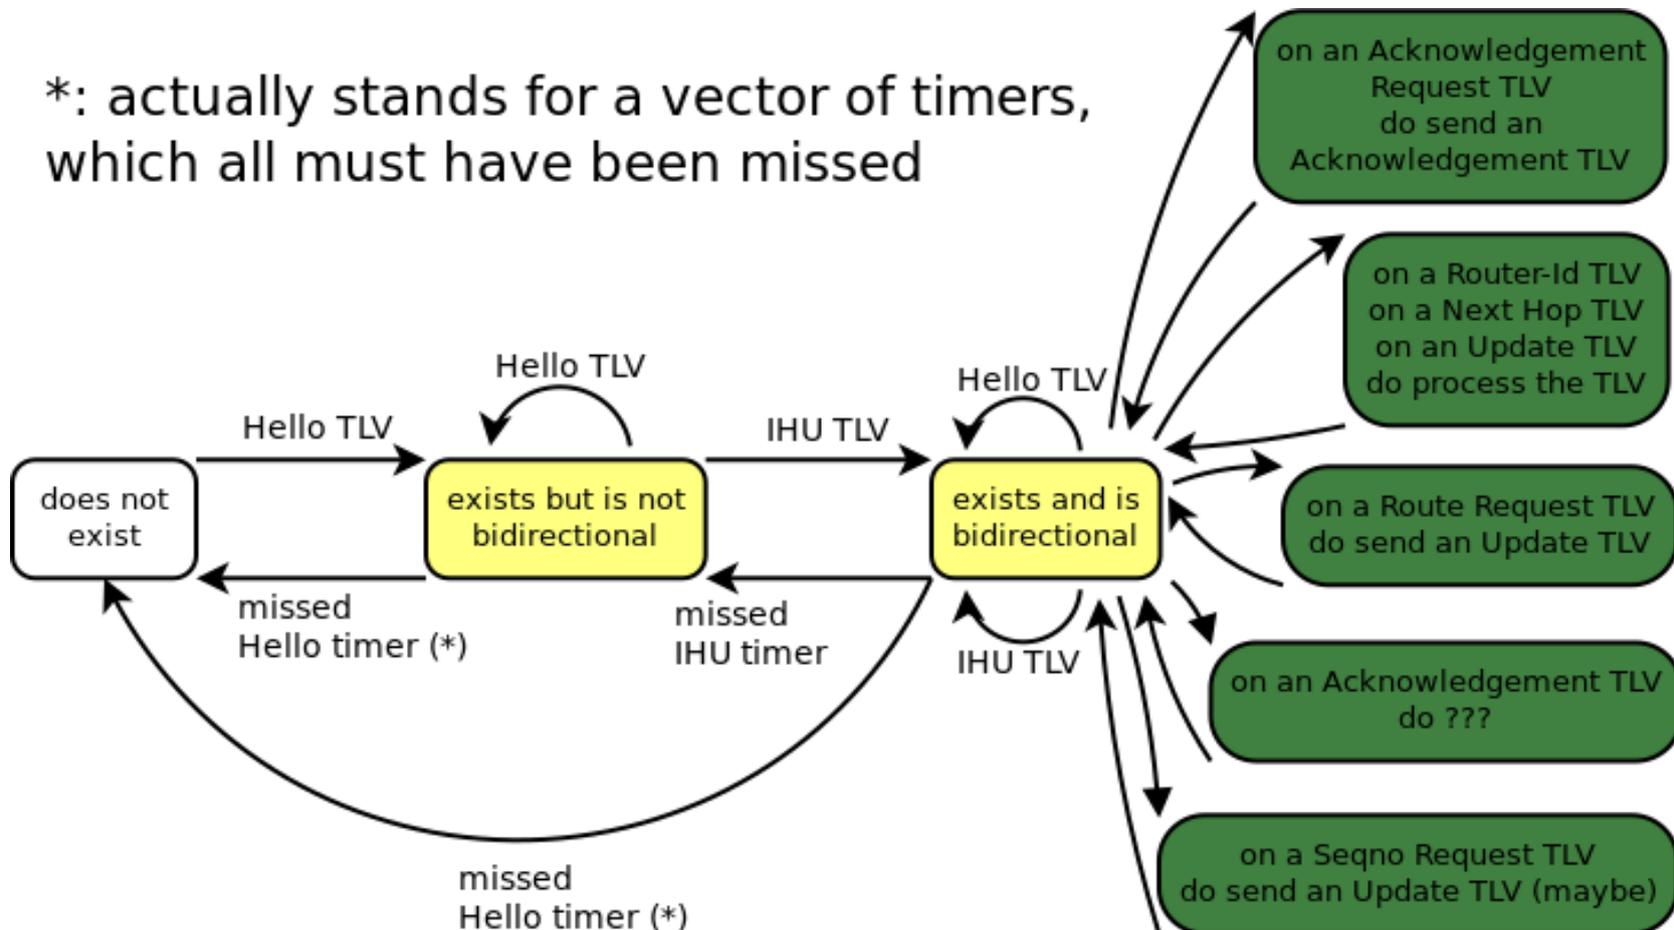


What it actually says

- A valid neighbour has to prove bidirectional reachability and the requirement works symmetrically both ways, **which is what was intended originally**.
- A valid neighbour has to keep sending both IHU TLVs and Hello TLVs, **which may be not quite precisely what was intended, though sufficiently close**. Anyway, it has proved to work and protocol properties coming out of this arrangement can be discussed.
- Rest of the TLVs doesn't depend on bidirectional reachability (as far as my interpretation of the spec goes), **which is likely not what was intended**. The next slide tries to address this with a slightly different FSM diagram.

Should it be like this?

*: actually stands for a vector of timers, which all must have been missed

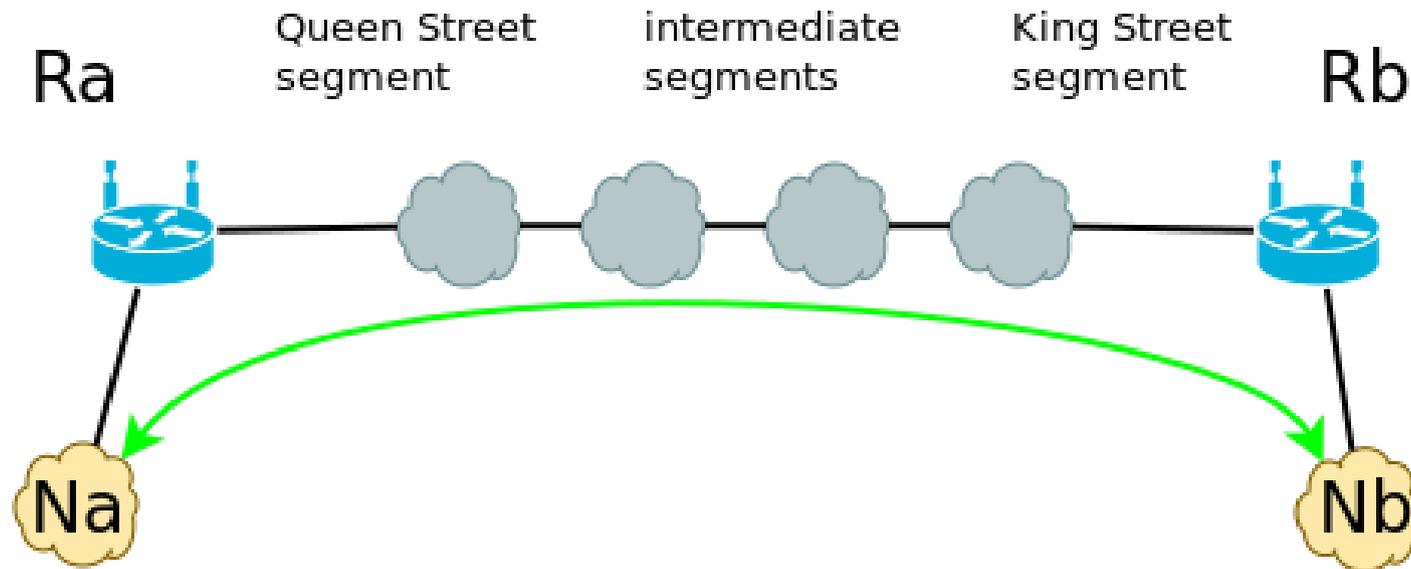


Discussion/decision points

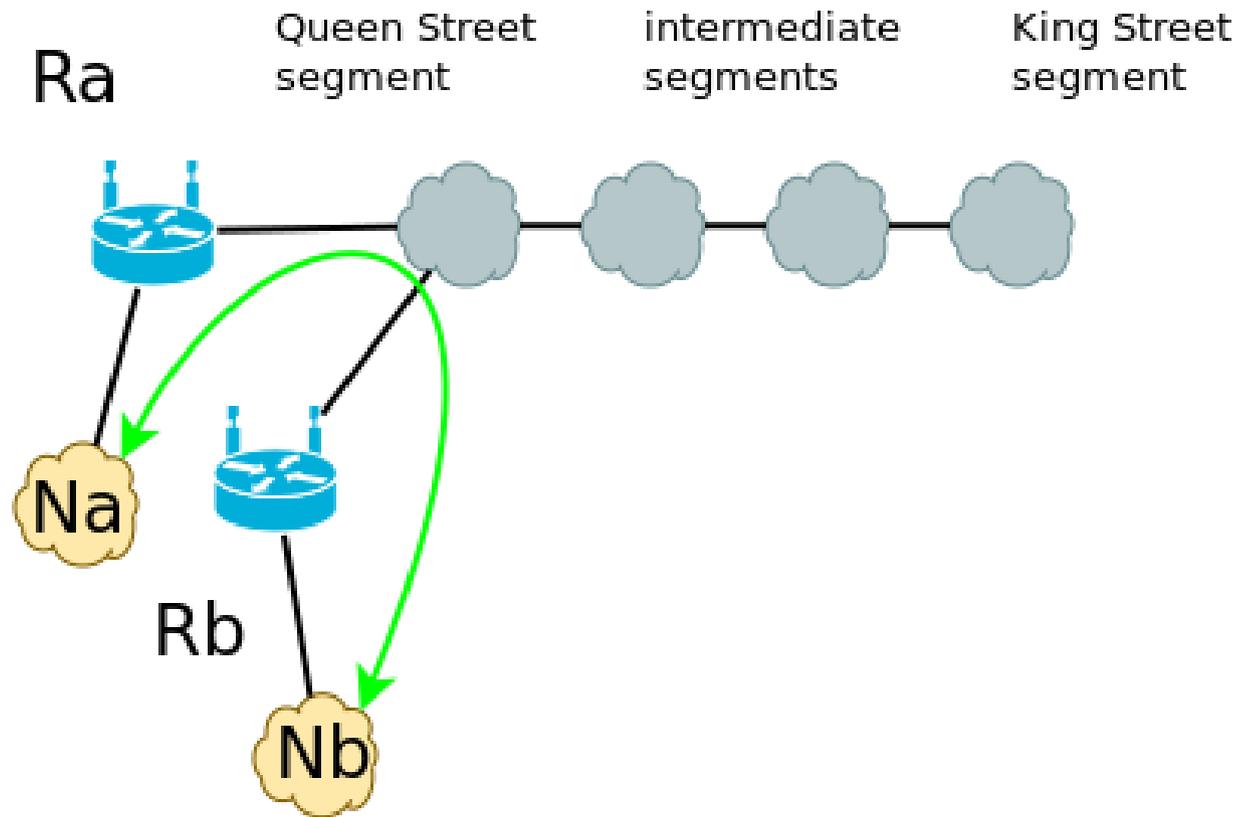
- Some of the thinking made in Appendix A.1 looks normative and could make Section 3 (Protocol Operation) more complete if moved there.
- The specification should explicitly state that the “green” TLVs must be ignored unless they come from a valid neighbour (or if some must not be, it should tell which TLVs and why).
- The dependency of a valid neighbour entry on receiving both Hello and IHU TLVs on time may not look obvious but is critical for a correct implementation. The specification should acknowledge or justify this detail, possibly in Section 3.4 (Neighbour Acquisition).
- Section 3.3 specifies clearly how to respond to an Acknowledgement Request TLV and discusses a bit when to send it but there is nothing defining how to process an expected/unexpected Acknowledgement (response) TLV or a lack of an expected response. This should be clarified but I have no suggestions how exactly.

II. Suggested updates to bidirectional reachability
detection to address a flaw in RFC 7298
authentication mechanism

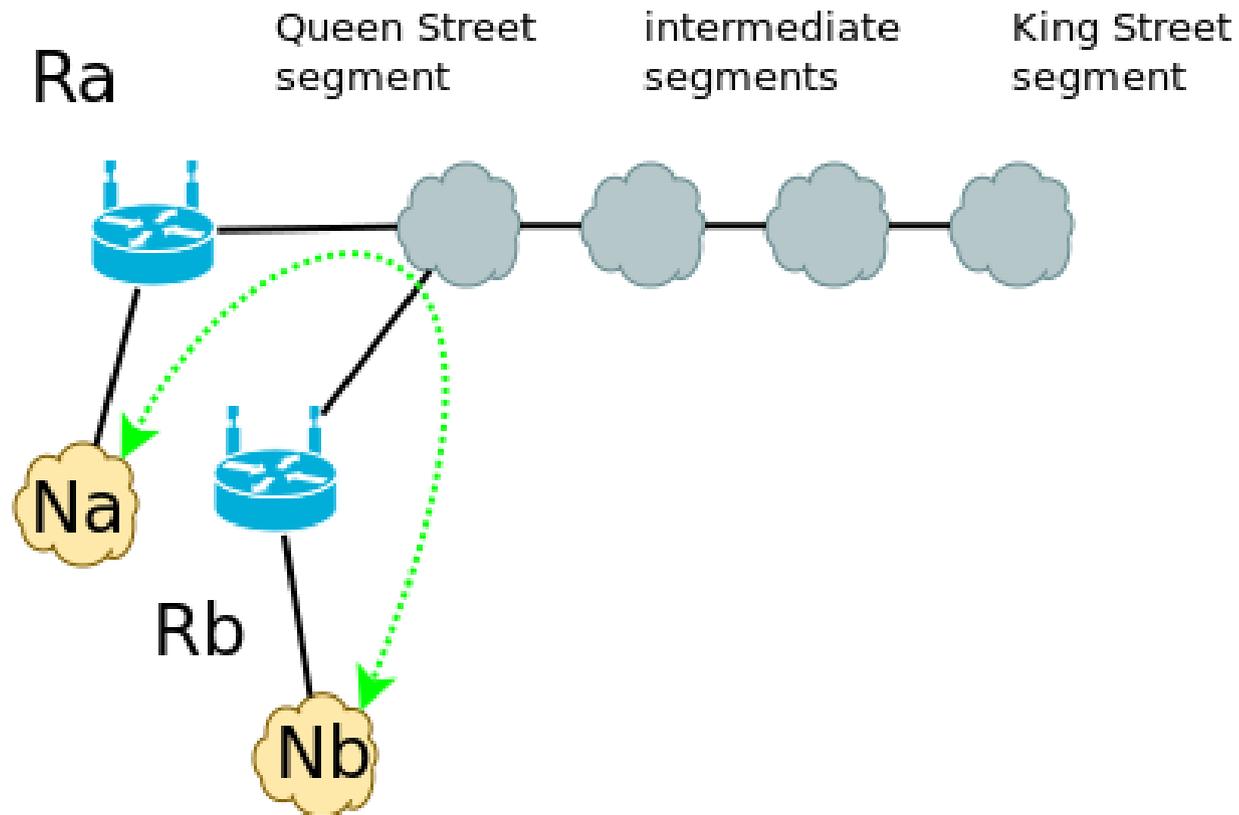
The attack (topology)



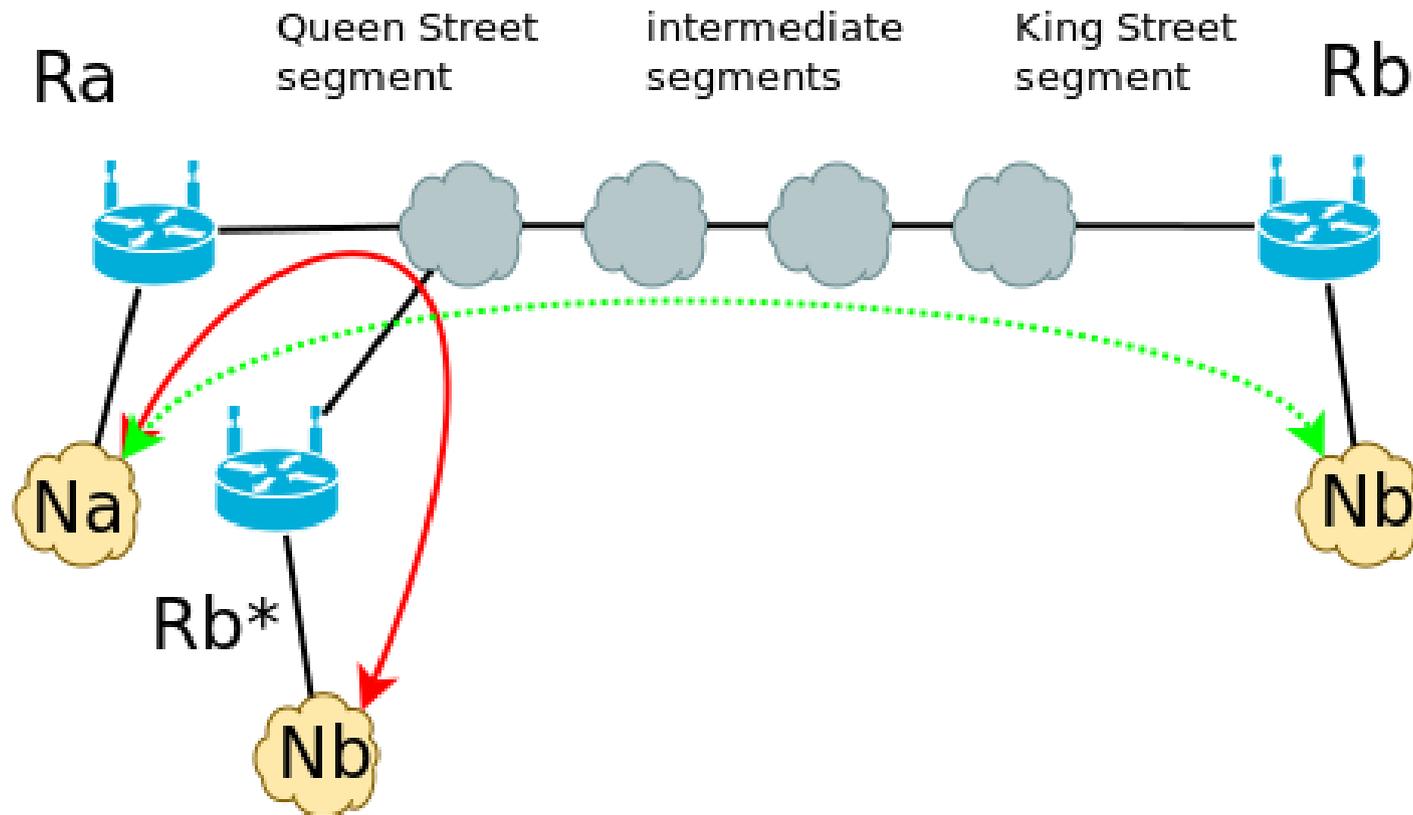
The attack (topology)



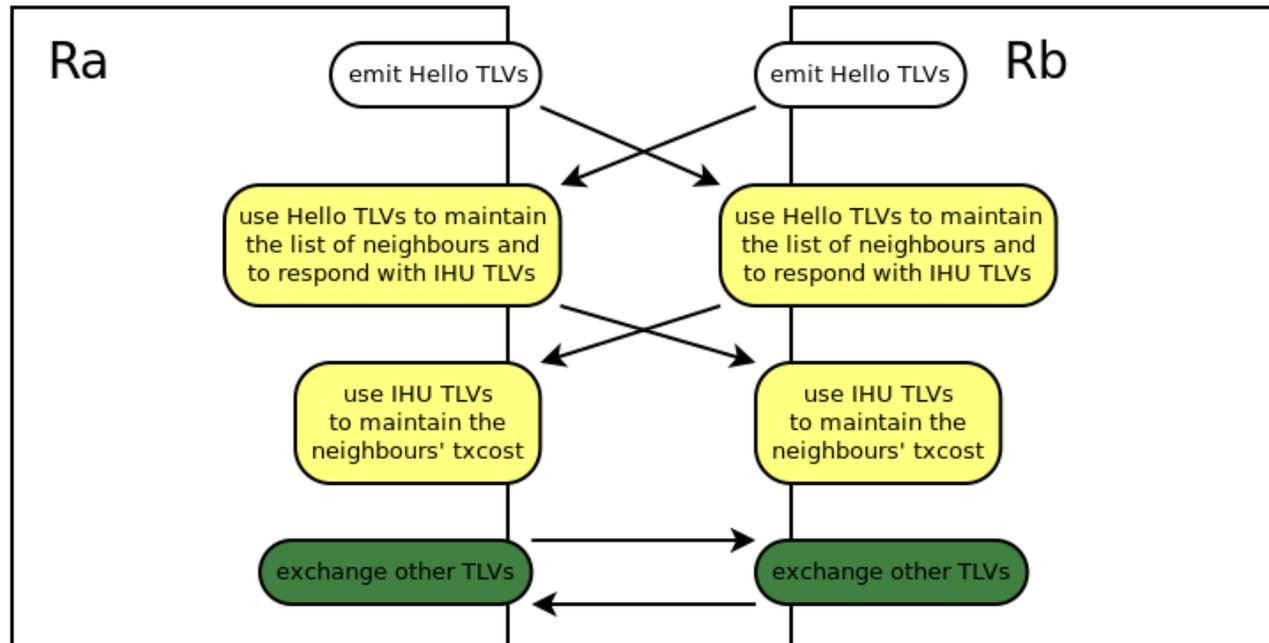
The attack (topology)



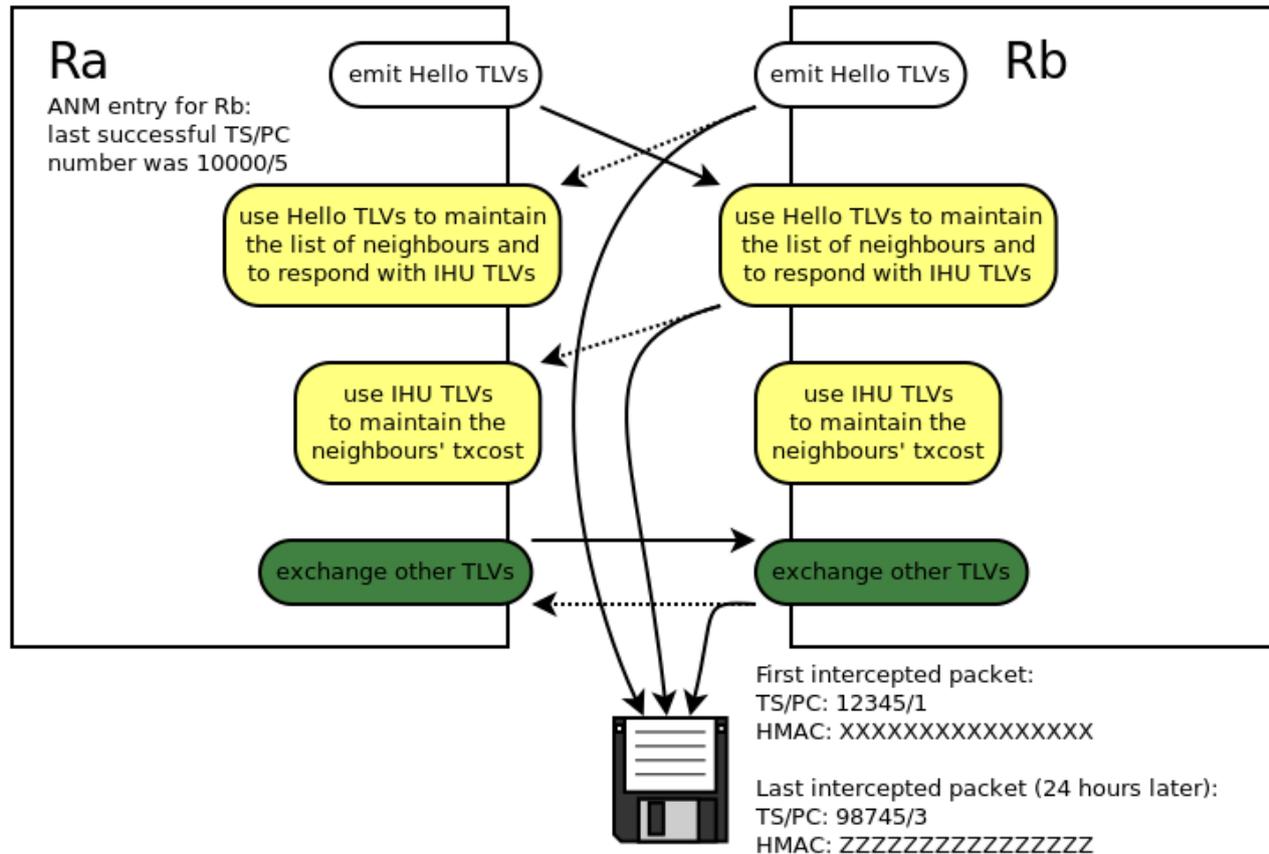
The attack (topology)



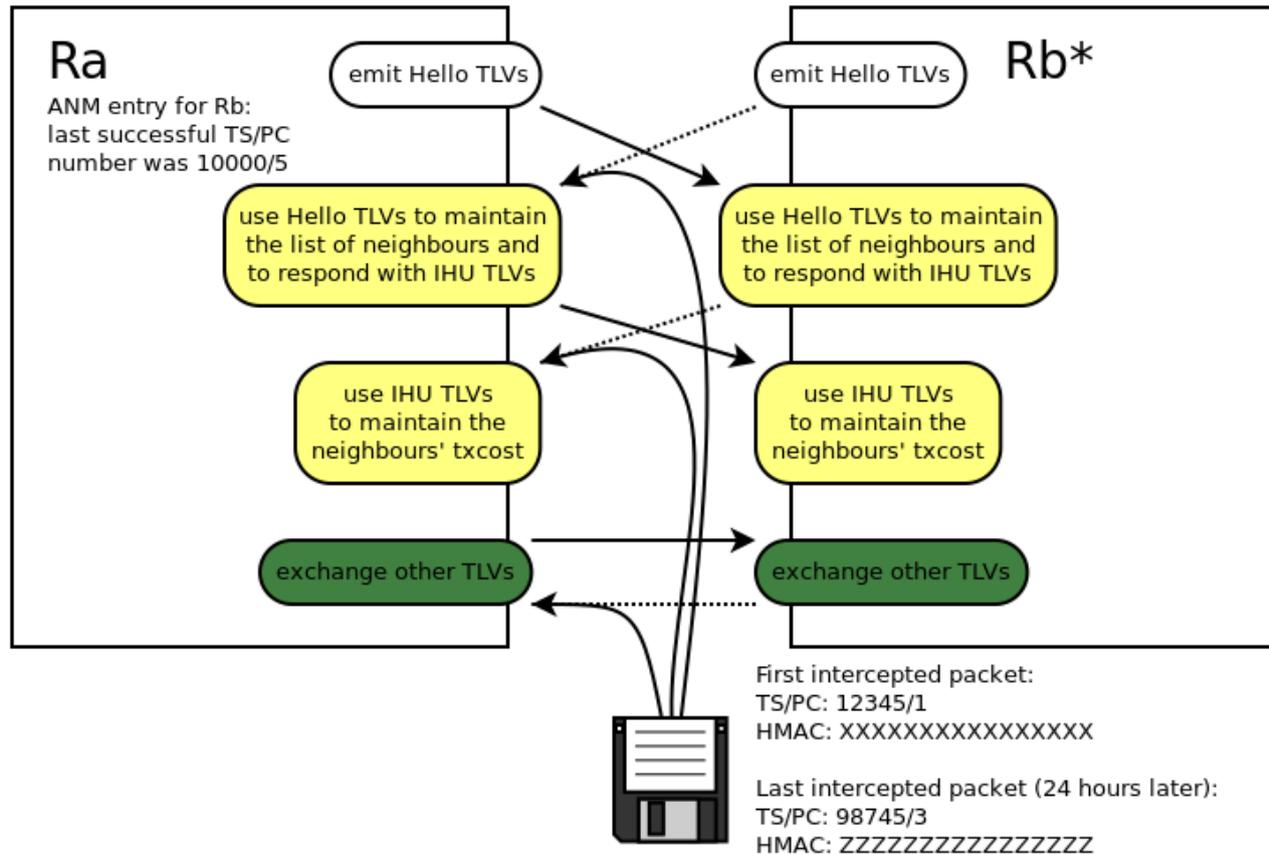
The attack (packet exchange)



The attack (packet exchange)



The attack (packet exchange)



The problem

The base protocol specification defines no means to relate a received IHU TLV to any of the recently sent Hello TLVs. The authentication mechanism design did not account for this property at the time and hence did not include its own measures to address it.

On the bright side is that now the time would be good to solve this in either the base protocol or the authentication mechanism or both.

Possible solution: basic principle

A very similar problem had been solved in TCP with the 3-way exchange and TCP sequence numbers. The same may work here: if one speaker sends a number together with its Hello TLV and then later the other speaker echoes the same number back together with its IHU TLV, then the first speaker can compare the received and the expected numbers to tell an out-of-date IHU TLV (or whole packet).

Possible solution: particular options

- Piggy-back a sub-TLV with the latest received Hello Seqno (16-bit) for respective neighbour on the IHU TLV.
- Idem, with Hello Seqno expanded to 32 or more bits through a change of TLV encoding or a sub-TLV.
- Idem, but leave Hello Seqno alone and use the TS/PC number (48-bit) for the same purpose.
- Let the existing RTT extension measure IHU age and compare it with a threshold.

Conclusions

- Both stated problems need to be addressed.
- The first one is simpler and hopefully requires only editorial work. If the WG has a consensus on any of the suggested changes right now, I can prepare and send proposed changes for the 6126-bis I-D to babel@ietf. Otherwise please state your alternatives.
- The second one requires a protocol design decision and feedback is welcome to make this decision well.

Thank you!