

Stateless Reverse Traceroute


Rolf Winter

<https://datatracker.ietf.org/doc/html/draft-heiwin-intarea-reverse-traceroute-stateless-03>

The 2 big questions before moving forward

- Encoding
 - The draft proposes to use ICMP for the signaling part (triggering traceroute probes and reporting results back)
 - Two main options:
 - Use Echo (ICMP codes 0/8 and ICMPv6 code 128/129)
 - Use Extended Echo (ICMP codes 42/43 and ICMPv6 codes 160/161)
- Stateless vs. stateful
 - Go fully stateless for the server
 - Go fully stateful for the server
 - Let the implementation decide

Recap: Reverse Traceroute



➡ What Traceroute reveals

➡ What Reverse Traceroute reveals

Question 1 - Encoding

Measurements - Echo (recap)

- Which ICMP type/code combinations will go through common middleboxes **today**? (Deployability)
- Echo with new code (12 middleboxes tested)

ICMP request	forwarded	filtered	bypassed
Type 8, code 1	11	1 ^{a)}	0
Type 8, code 2	11	1 ^{a)}	0
Type 7, code 0	1	7	4
Type 252, code 0	1	6	5

^{a)} Response dropped

- **Verdict:** Echo with new code **works** through most NAT implementations today

Measurements - Extended Echo

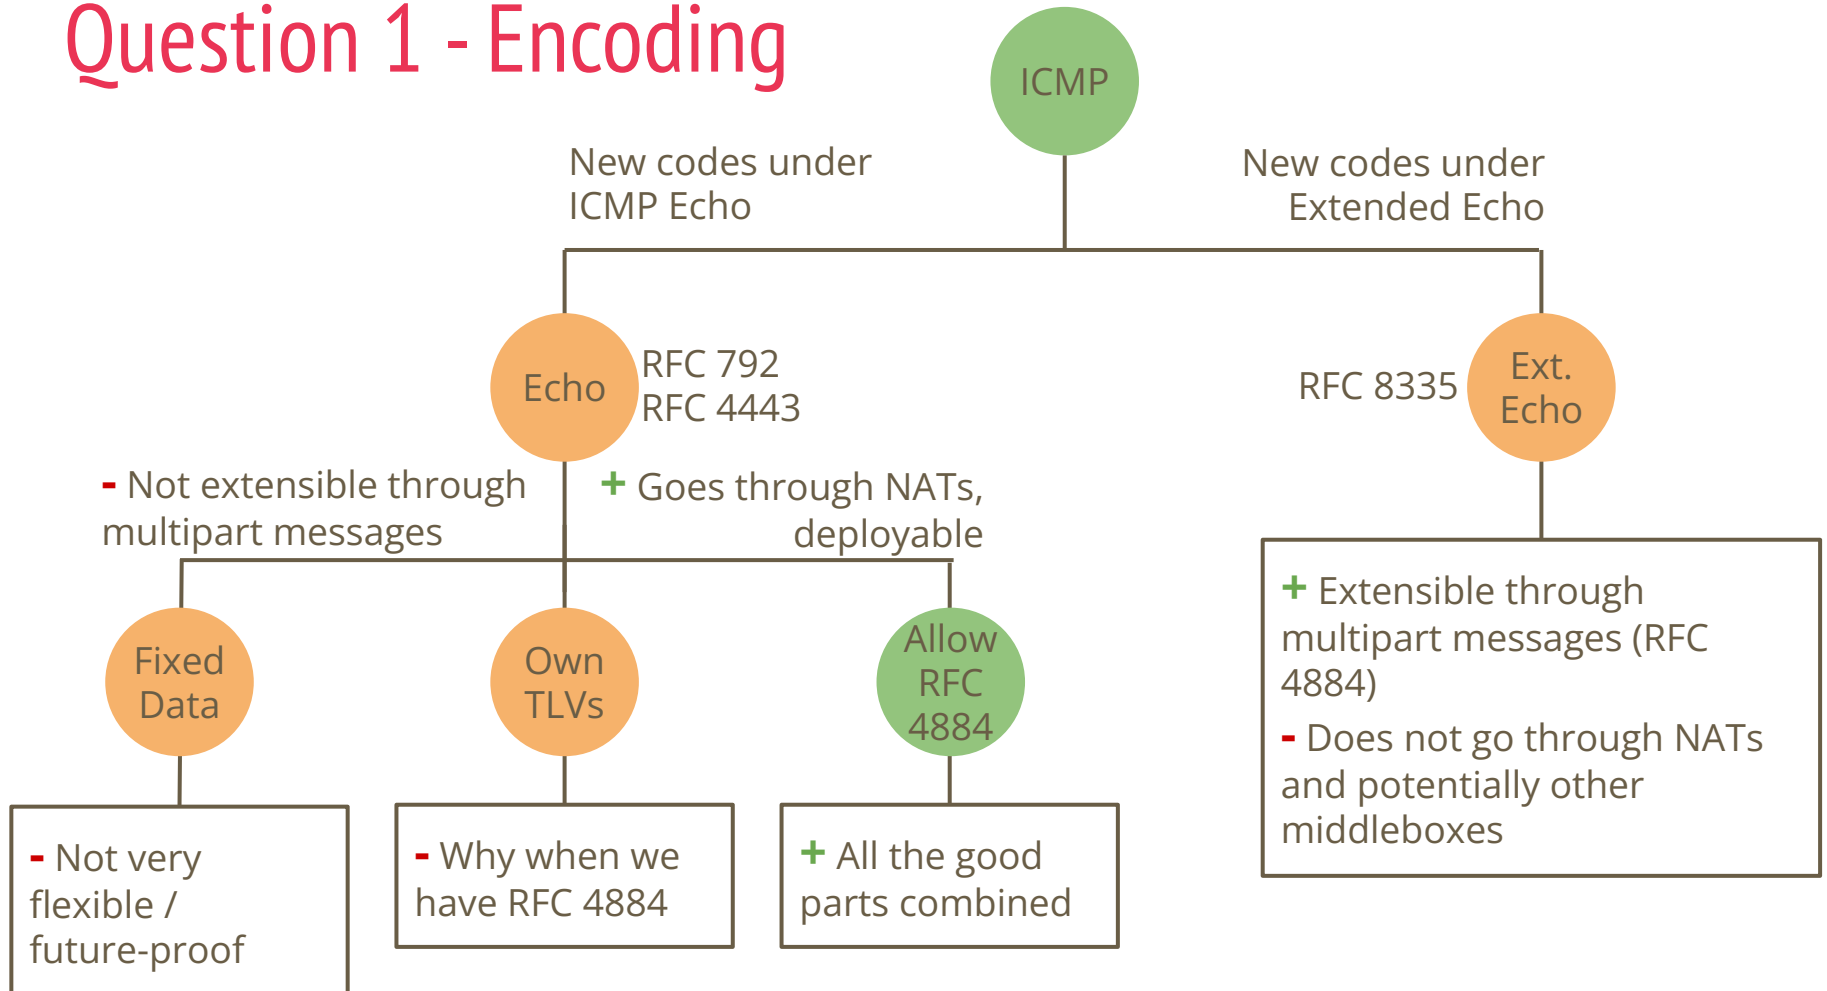
- Will Extended Echo go through common middleboxes **today?** (Deployability)
- Extended Echo (23 middleboxes tested)
 - Careful: Extended Echo was standardized only 2018
 - Only 6 boxes started sales after 2018 (none of which support Extended Echo)

ICMP request	forwarded	Request forwarded & response filtered	filtered	Request bypassed	
				Response filtered	Response delivered (but also without address translation)
Type 42, code 0	1 ^{a)}	1	13	5	3

^{a)} 2014 model

- **Verdict:** Extended Echo will way more often than not **not work** today

Question 1 - Encoding



Question 2 - Stateful vs. stateless

For stateless, we need space in (ICMP Time Exceeded) packets

- Based on RIPE Atlas data, we only have enough space in 50% of the packets
 - That means, we can find out who is on the path back to the requester, but no RTT estimation or anything fancy

Response size [bytes]	8	12	20 - 28	32 - 39	44	>=48
% of responses	48.6%	0.1%	0.4%	0.5%	0.4%	50%

Question 2 - Stateless vs. Stateful

- Stateless vs. stateful refers to whether the server has to keep state, while a traceroute probe is on its way
- Stateless
 - + No state maintenance needed
 - - RTT estimations are not always possible - measurements suggest 50% of the times
 - - no fancy functions possible, that would require state, such as tracerouting somewhere else (not back to the requester)
- Stateful
 - + RTT estimations always work
 - + Fancy functions possible
 - - need to keep and maintain a little bit of state
- **Suggestion:** leave it to the implementation

Next steps

- Re-submit the draft with what we decide
- Then ask for adoption

- As a reminder, the questions are:
- Encoding (4 options)
 - ICMP Echo
 - Fixed data structure
 - Hand-rolled TLVs
 - **Allow multipart with the newly defined codes**
 - ICMP Extended Echo
- Stateful vs. Stateless (3 options)
 - Always stateless
 - Always stateful
 - **Let the implementation decide**