ATR: Additional Truncation Response for Large DNS Response

draft-song-atr-large-resp-01

Davey Song (speaker)/BII lab,

2018-07-17 @IETF 102 Montreal

天地互连

Document history

- draft-song-atr-large-resp-00
  - An simple improvement on authoritative server by replying additional truncated response just after the normal large response
  - Experimental document
- draft-song-atr-large-resp-01
  - An ATR indicator is introduced with a AT bit in EDNS0 OPT header
  - More operational considerations on ATR timer, ATR payload size and Less-aggresiveness of ATR
  - Standards Track document

Background

- Two orthogonal ways handling large DNS response
  - Fallback to TCP via TrunCation bit
  - Use EDNS(0) generate larger response avoid TCP fallback
- More public evidence and concerns on packets drop caused by IPv6 fragmentation in DNS
  - RFC7872 reports more than 30% drop rates for sending fragmented packets
  - An APNIC measurement report says more than 37% of endpoints using IPv6-capable DNS resolver can not receive a fragmented IPv6 response over UDP

ATR in one slides

- Decouple TCP fallback and EDNS(0) and make them paralleled
- ATR adds an additional response packets to “trail” a fragmented UDP response

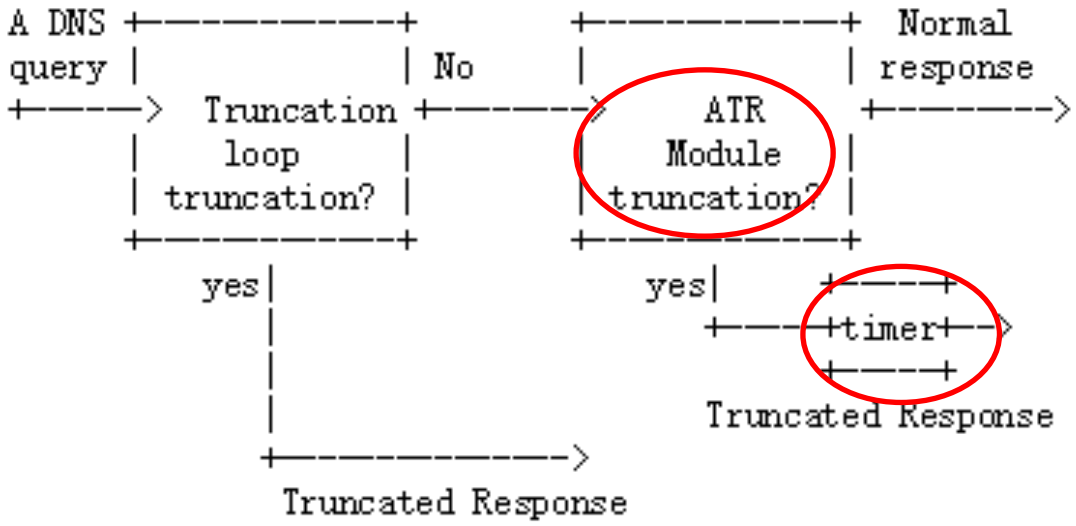

ATR payload size and ATR timer is two operational parameters

天地互连

The Intention of ATR

Today:

- If the client cannot receive large truncated responses then it will need to timeout from the original query,
- Then re-query using more resolvers,
- Timeout on these queries
- Then re-query using a 512 octet EDNS(0) UDP buffersize
- Then get a truncated response
- Then re-query using TCP

Geoff Huston , IEPG Meeting - March 2018 @ IETF 101
Source: http://iepg.org/2018-03-18-ietf101/geoff.pdf

The Intention of ATR

ATR

Today:

- If the client cannot receive large truncated responses then it will need to timeout from the original query,
- Then re-query using more resolvers,
- Timeout on these queries
- Then requery using a 512 octet EDNS(0) UDP buffersize
- Then get a truncated response within a few ms
- Then requery using TCP

Geoff Huston , IEPG Meeting - March 2018 @ IETF 101

Source: http://iepg.org/2018-03-18-ietf101/geoff.pdf

The Operation of ATR

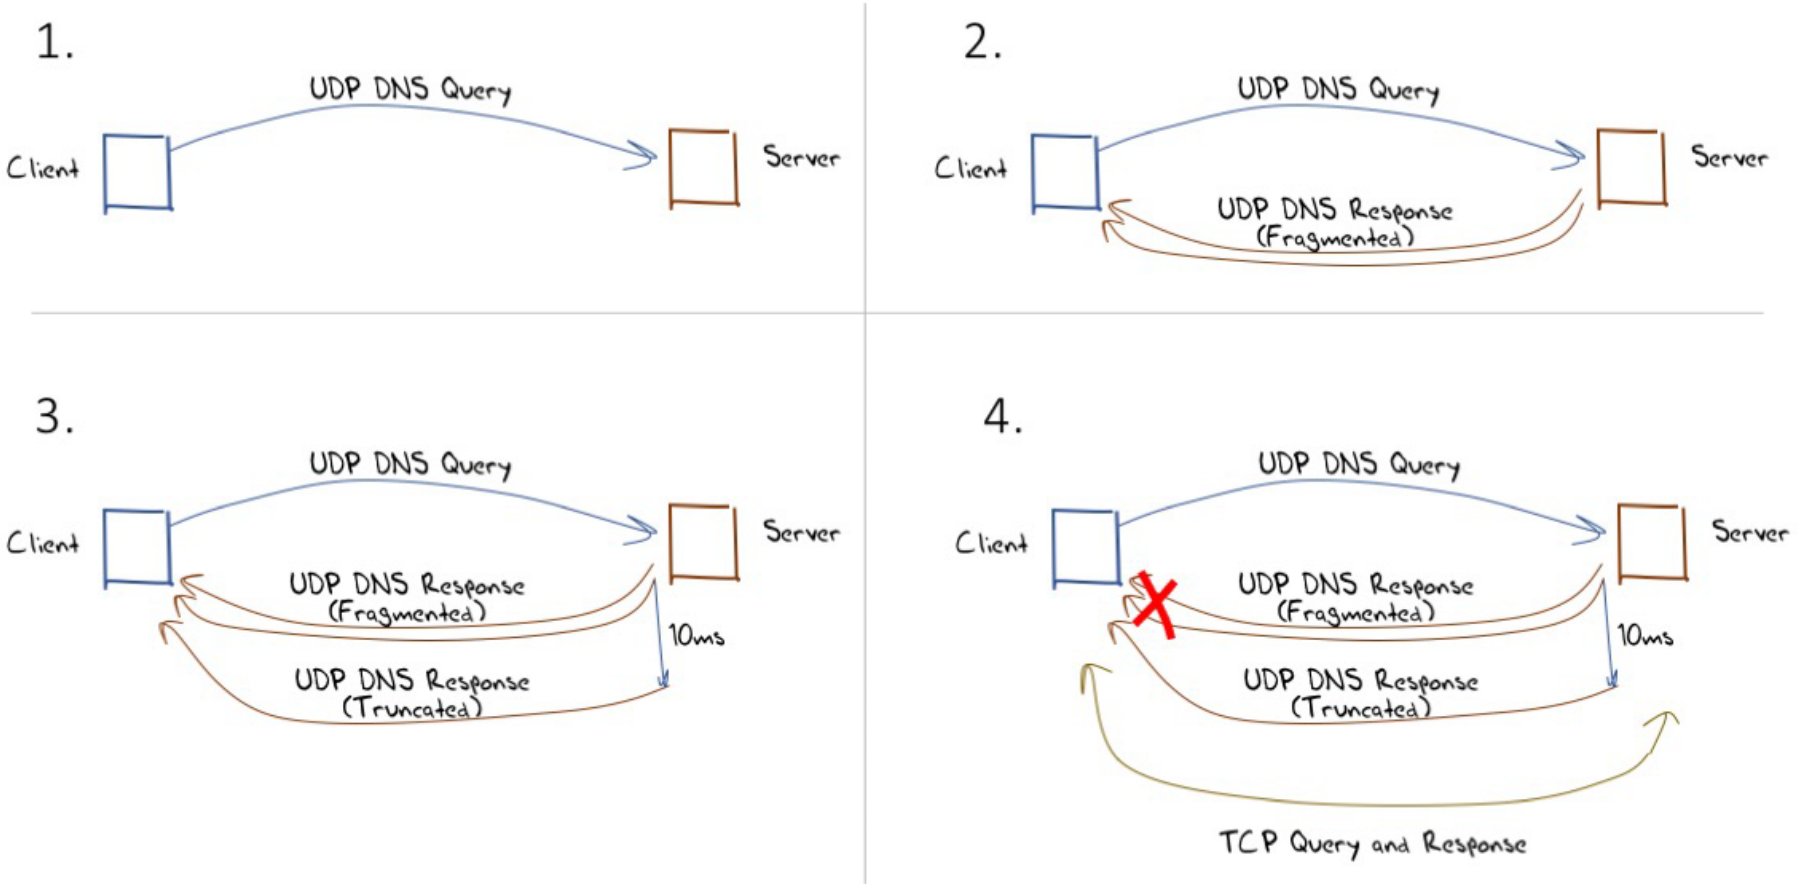

"How well does ATR actually work?" Geoff Huston , Apr 2018

Source :https://blog.apnic.net/2018/04/16/how-well-does-atr-actually-work/

天地互连

The benefit of ATR

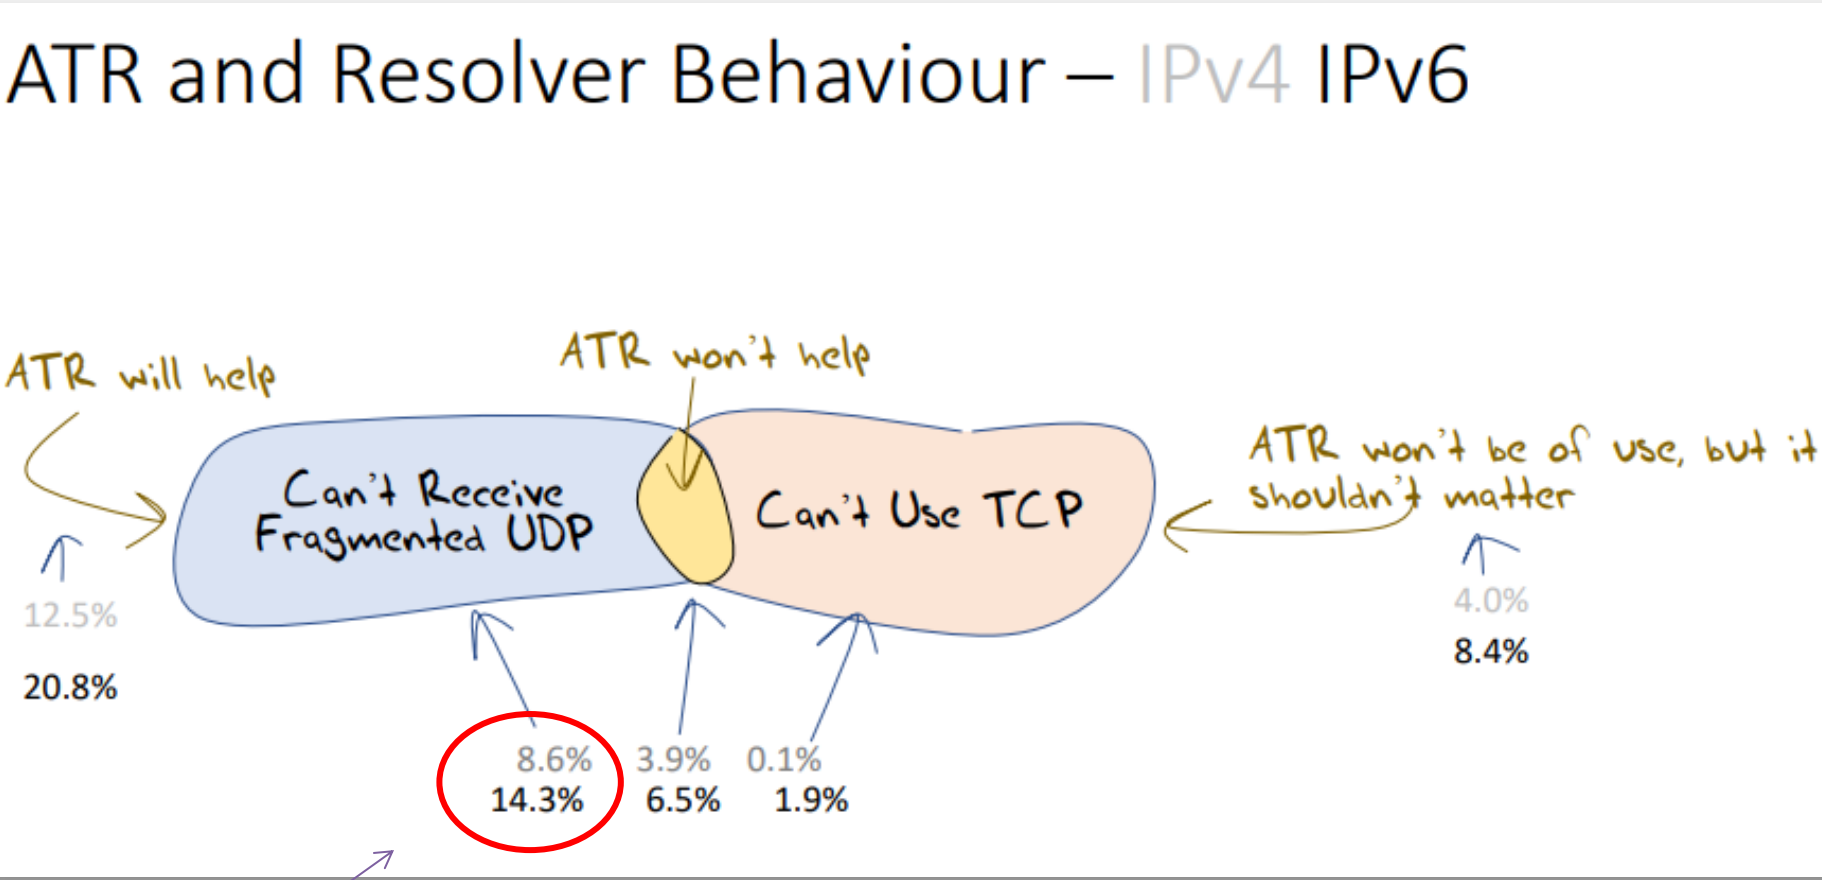

ATR benefits 68.6 IPv4 and 68.7 % IPv6 affected users

Source: http://iepg.org/2018-03-18-ietf101/geoff.pdf

天地互连

Intention of ATR indicator

- Intention to distinguish these ATR responses in a specific way from an ordinary truncated response
- Enable people to log cases where these ATR responses were received without having already received a (reassembled) UDP response to the query
- Indications to flag problematic name servers where people should restrict maximum EDNS to a lower value than the default 4096 that currently use

AT flag bit in EDNS0 OPT Header

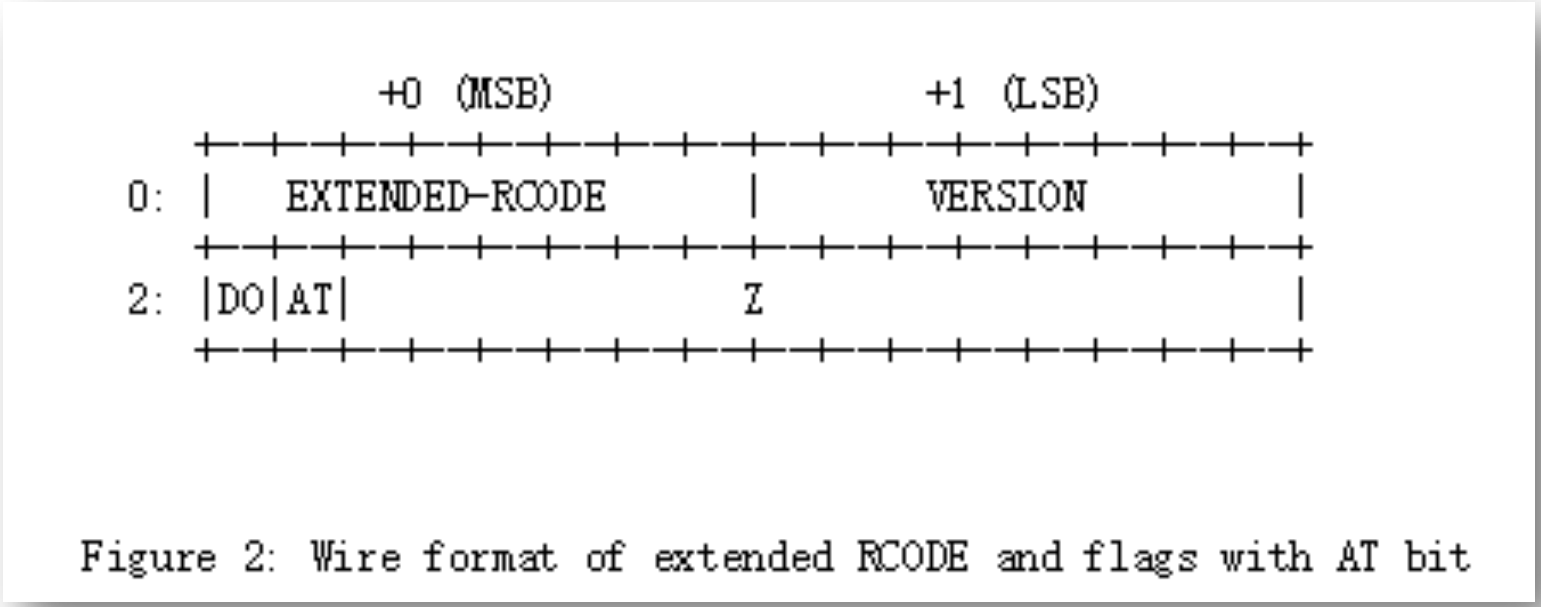

Figure 2: Wire format of extended RCODE and flags with AT bit

- Setting the AT bit to one in a response indicates to the resolver that the response is an ATR response.
- The AT bit cleared(set to zero) indicates the response is a ordinary response.

Source : https://tools.ietf.org/html/draft-song-atr-large-resp-01

天地互连

Operational Considerations

- ATR timer
  - To avoid the impact of network reordering(RO)
  - Less than 50 ms for large site DNS
- ATR payload size
  - 1472 octets for IPv4, 1232 octets for IPv6
- Less aggressiveness of ATR
  - ATR may respond TC=1 responses at a low possibility, such as 10%.
  - Reply ATR response selectively (identify cases ATR ignored)
- Implement a separate daemon of ATR without modify authoritative server

Next step

- Do you like ATR?
- Do you like ATR indicator , AT bit?
- Adoption in DNSOP?
- Standards track or experimental?

天地互连