Attacking the Network Time Protocol (NTP)

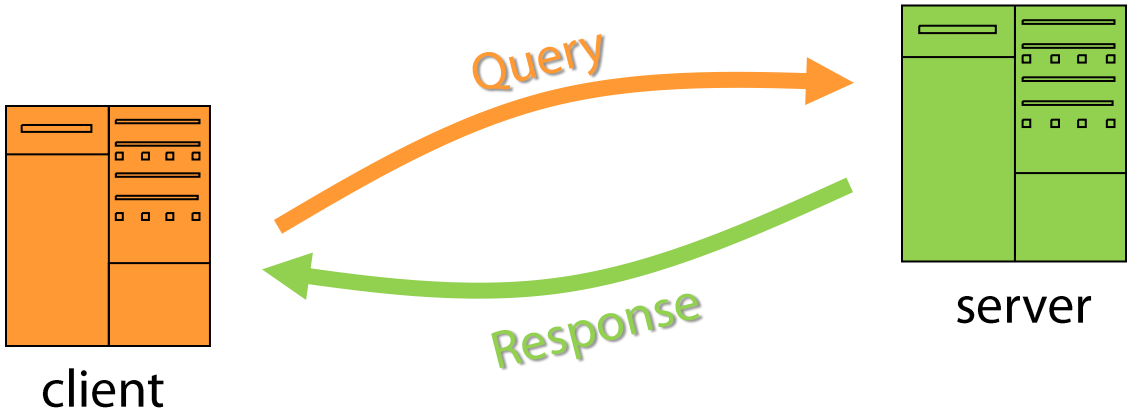

BOSTON UNIVERSITY

Aanchal Malhotra
Isaac E. Cohen, Erik Brakke
Sharon Goldberg

outline of the talk

- **Background**
  - How does NTP work?
  - How does NTP client take time?

- **Our client/server mode attacks**
  - Denial of Service by Spoofed Kiss-of-Death (off-path)
  - Denial of Service by Priming the Pump (off-path)
  - Timeshifting by IPv4 Packet Fragmentation (off-path)

- **Broadcast mode attacks**
- **Other attacks (if time)**

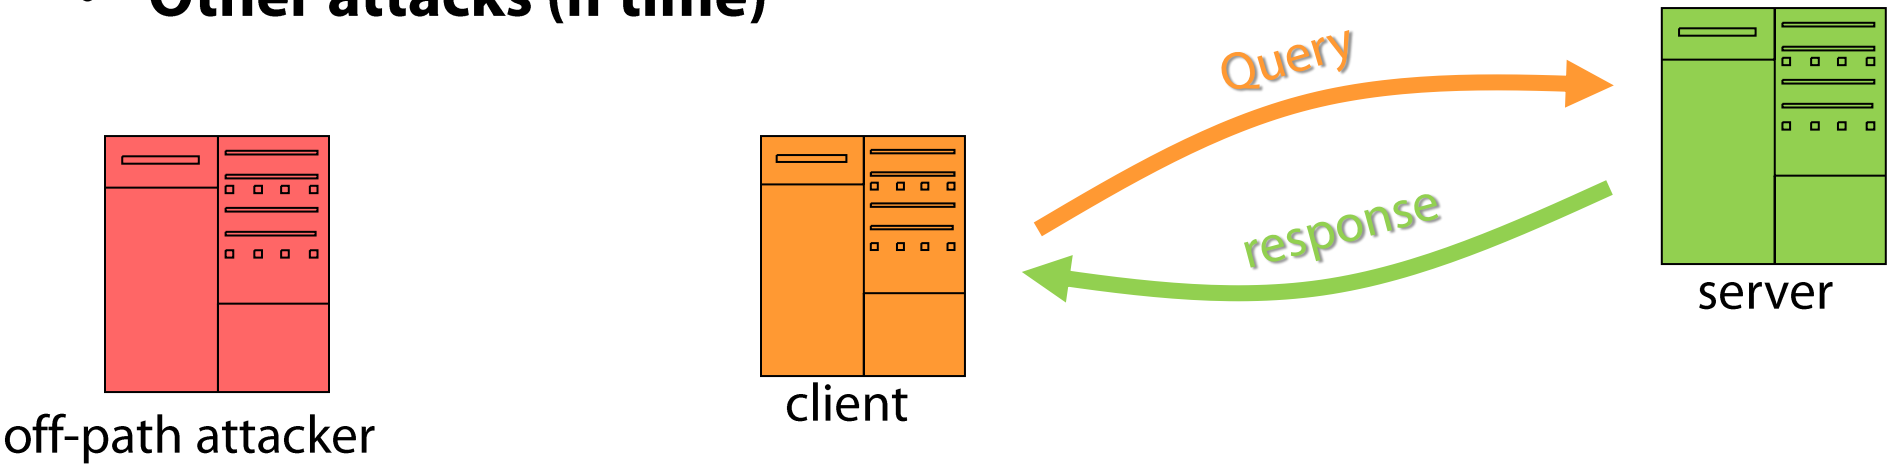

background: how does NTP work?

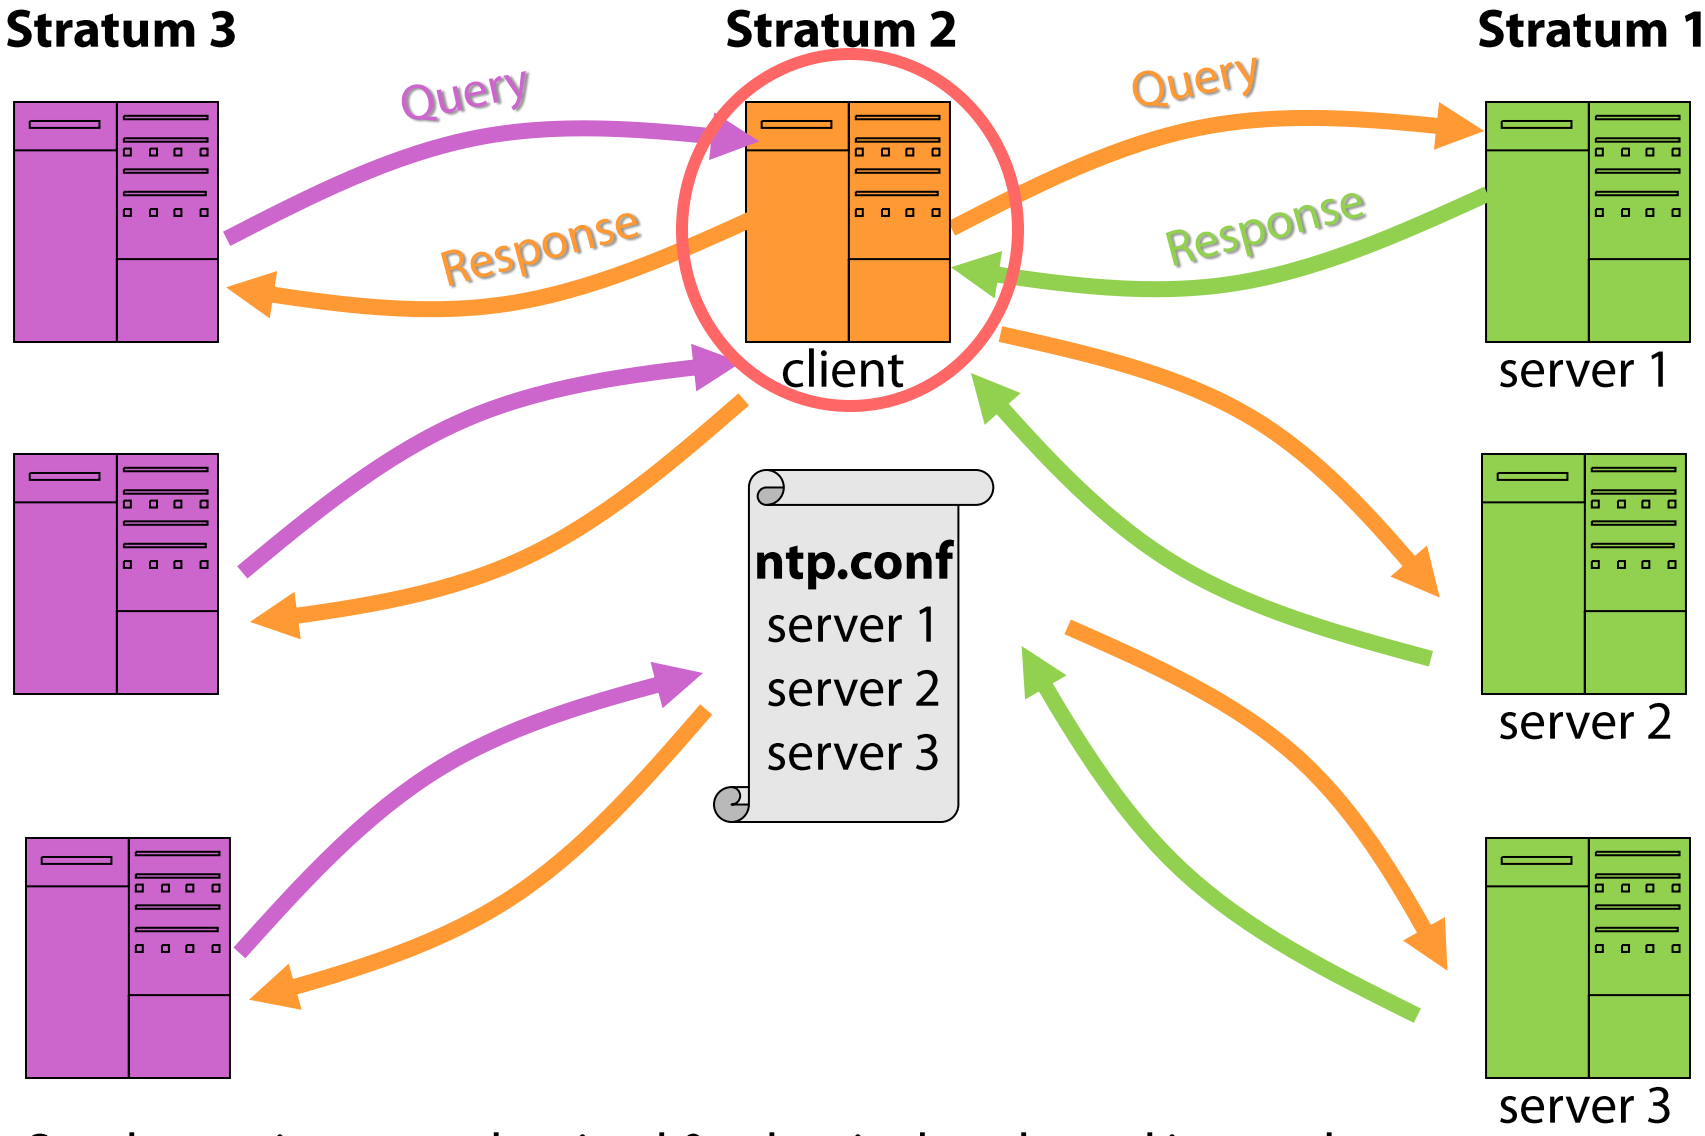

- Sends queries at randomized & adaptively-selected intervals
- Requires certain number of self-consistent responses to update its clock

background: how does NTP work?

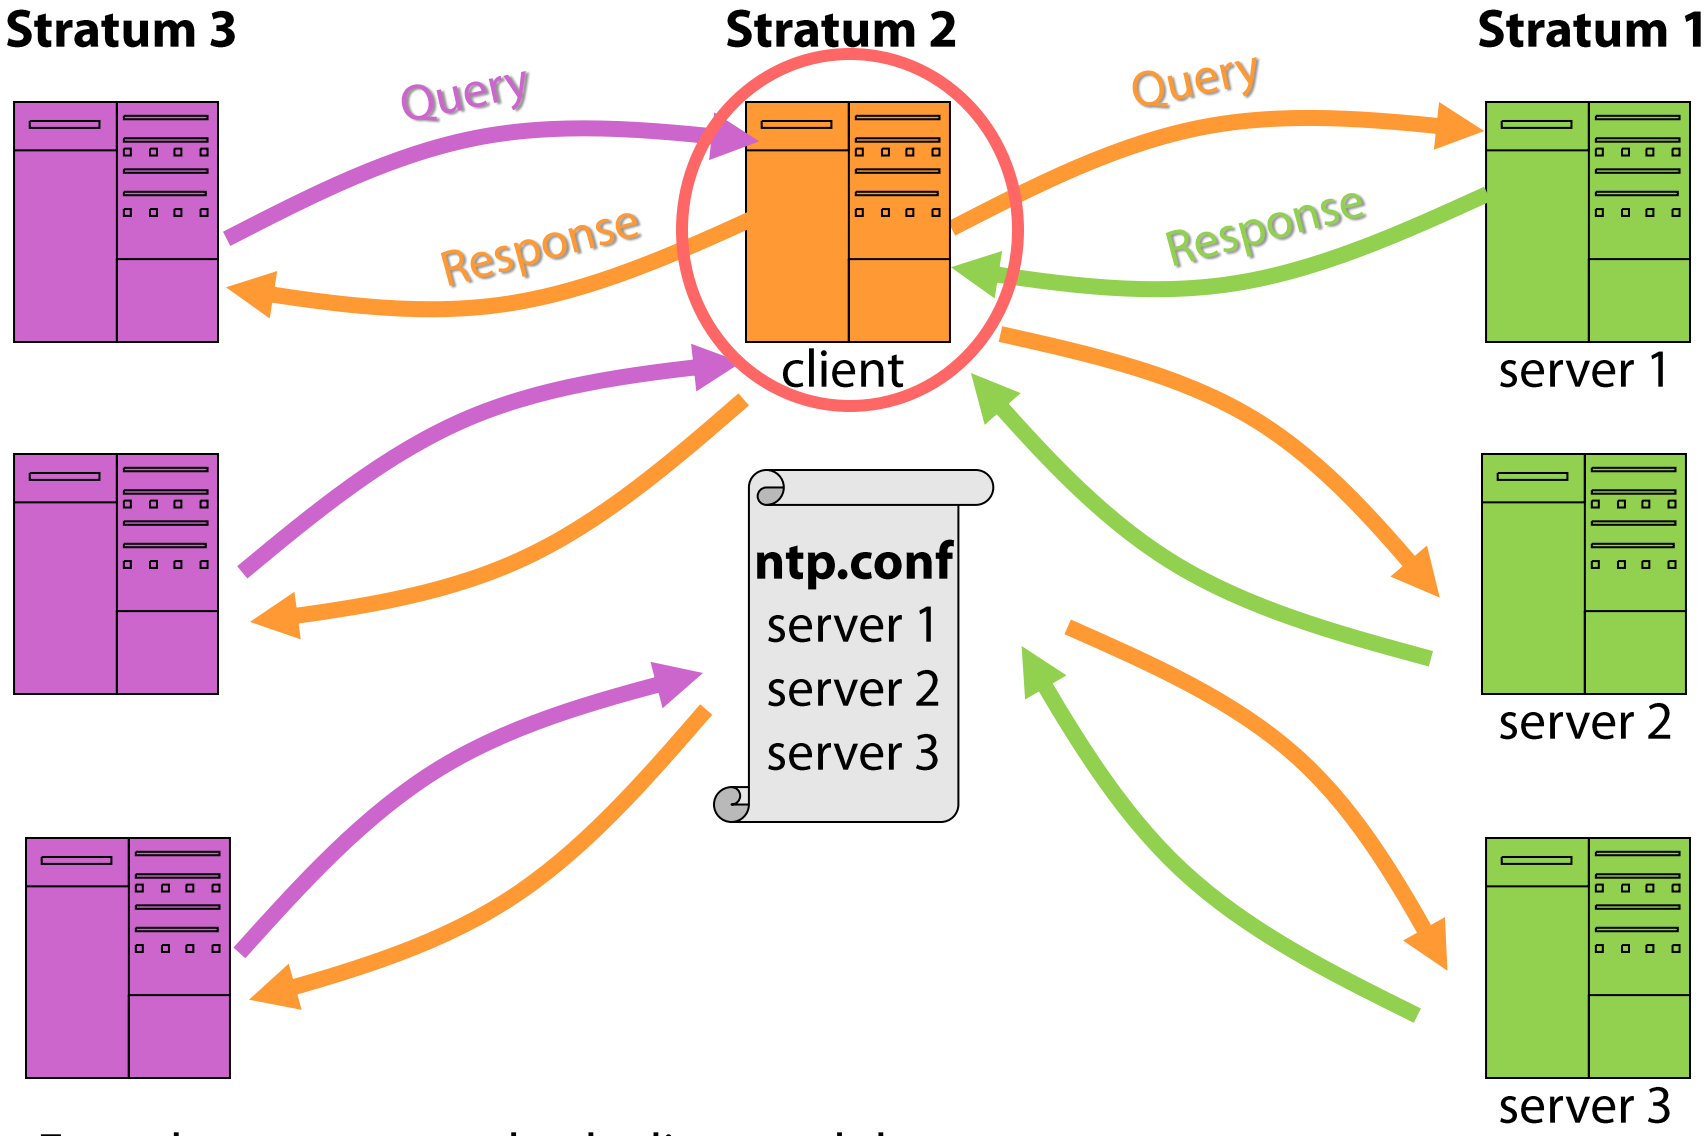

- Every host can act as both client and the server
- My laptop will answer queries from public Internet

the state of crypto in NTP

NTP's crypto is rarely used in practice

- Symmetric crypto
  - Uses *MD5(key||message)* [RFC 5905] (insecure!)
  - No in-band mechanism for key distribution
- Asymmetric crypto
  - Autokey Protocol [RFC 5906] is not a standards-track doc
  - Crypto is badly broken [S. Röttger' 2012]

Our zmap scan (July 2016) found 3.9M IPs revealing NTP crypto state

- Only 78K systems have all associations authenticated (2%)

IETF: Lots of activity lately in IETF to develop a secure NTP

- NTS (Network Time Security)
- Very fluid right now, but potentially based on DTLS.

We attack the NTPv4 spec [RFC 5905]

and its reference implementation (ntpd v4.2.8p2 & ntpd v4.2.6p5)

We assume NTP messages are not cryptographically authenticated.

non-crypto authentication with origin timestamp (T_1)

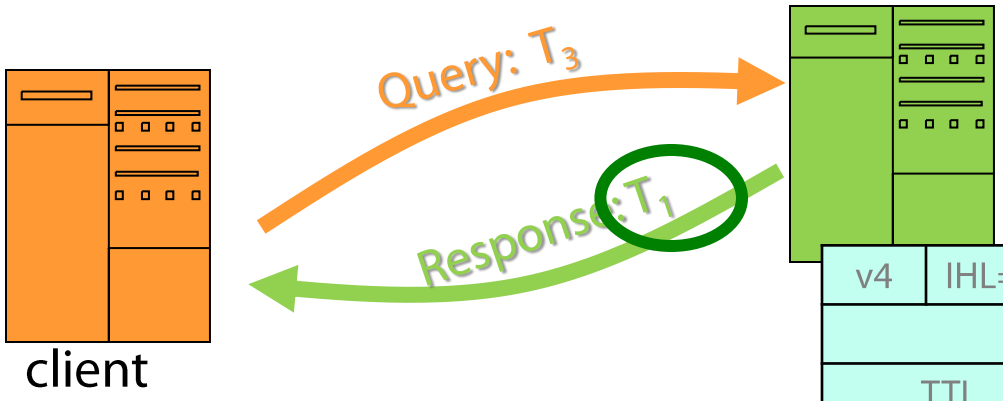

Analogous to
- UDP source port randomization
- TCP sequence no randomization

TEST2: Match
T3 in Query to **T1 in Response**.

v4	IHL=20	TOS	Total length = 76		
IPID			x	DF	MF	Frag Offset
TTL	Protocol = 17	IP Header Checksum			
Source IP					
Destination IP					
Source Port = 123			Destination Port = 123		
Length = 76			UDP Checksum		
LI	v4	Response	Stratum	Poll	**Precision**
Root Delay					
Root Dispersion					
Reference ID					
Reference Timestamp					
T_1 = Origin Timestamp					
T_2 = Receive Timestamp					
T_3 = Transmit Timestamp					

seconds
fractional seconds

64 bits

*ntpd does not randomize source port!

non-crypto authentication with origin timestamp (T_1)

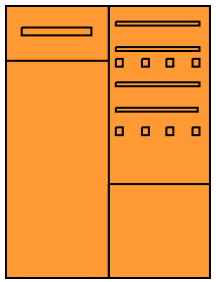

client

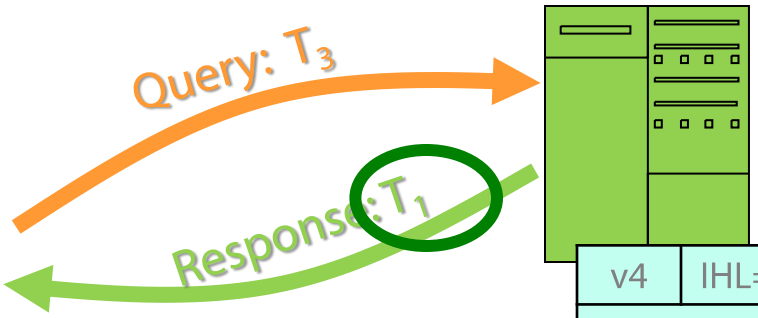

Analogous to
- UDP source port randomization
- TCP sequence no randomization

TEST2: Match
T3 in Query to **T1 in Response**.

v4	IHL=20	TOS	Total length = 76
IPID		x DF MF	Frag Offset
TTL	Protocol = 17		IP Header Checksum
Source IP			
Destination IP			
Source Port = 123			Destination Port = 123
Length = 76			UDP Checksum
LI v4	Response	Stratum	Poll / **Precision**
Root Delay			
Root Dispersion			
Reference ID			
Reference Timestamp			
T_1 = Origin Timestamp			
T_2 = Receive Timestamp			
T_3 = Transmit Timestamp			

seconds	
precision	**randomness**

64 bits

*ntpd does not randomize source port!

outline of the talk

- Background
  - How does NTP work?
  - How does NTP client take time?

- Our client/server mode attacks
  - **Denial of Service** by Spoofed Kiss-of-Death (off-path)
  - **Denial of Service** by Priming the Pump (off-path)
  - Timeshifting by IPv4 Packet Fragmentation (off-path)

- Broadcast mode attacks

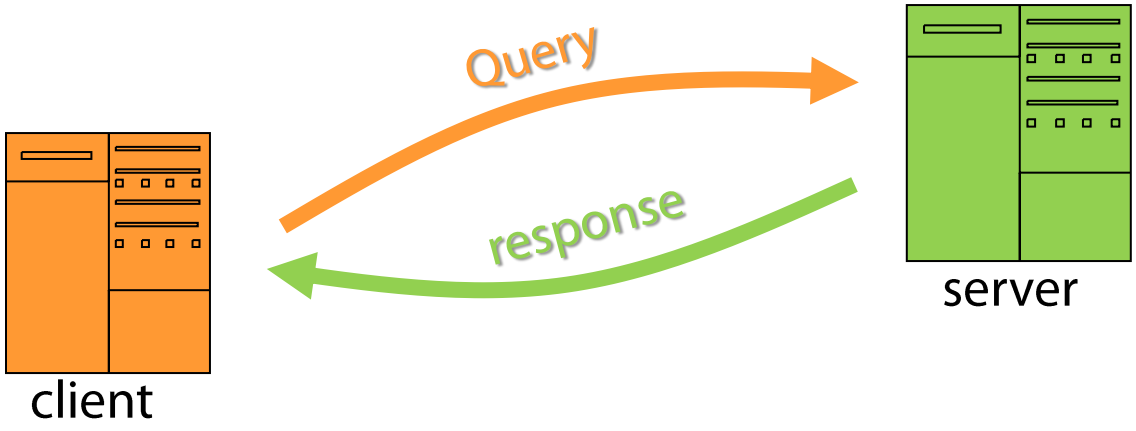

denial of service via spoofed kiss-o-death (KoD)

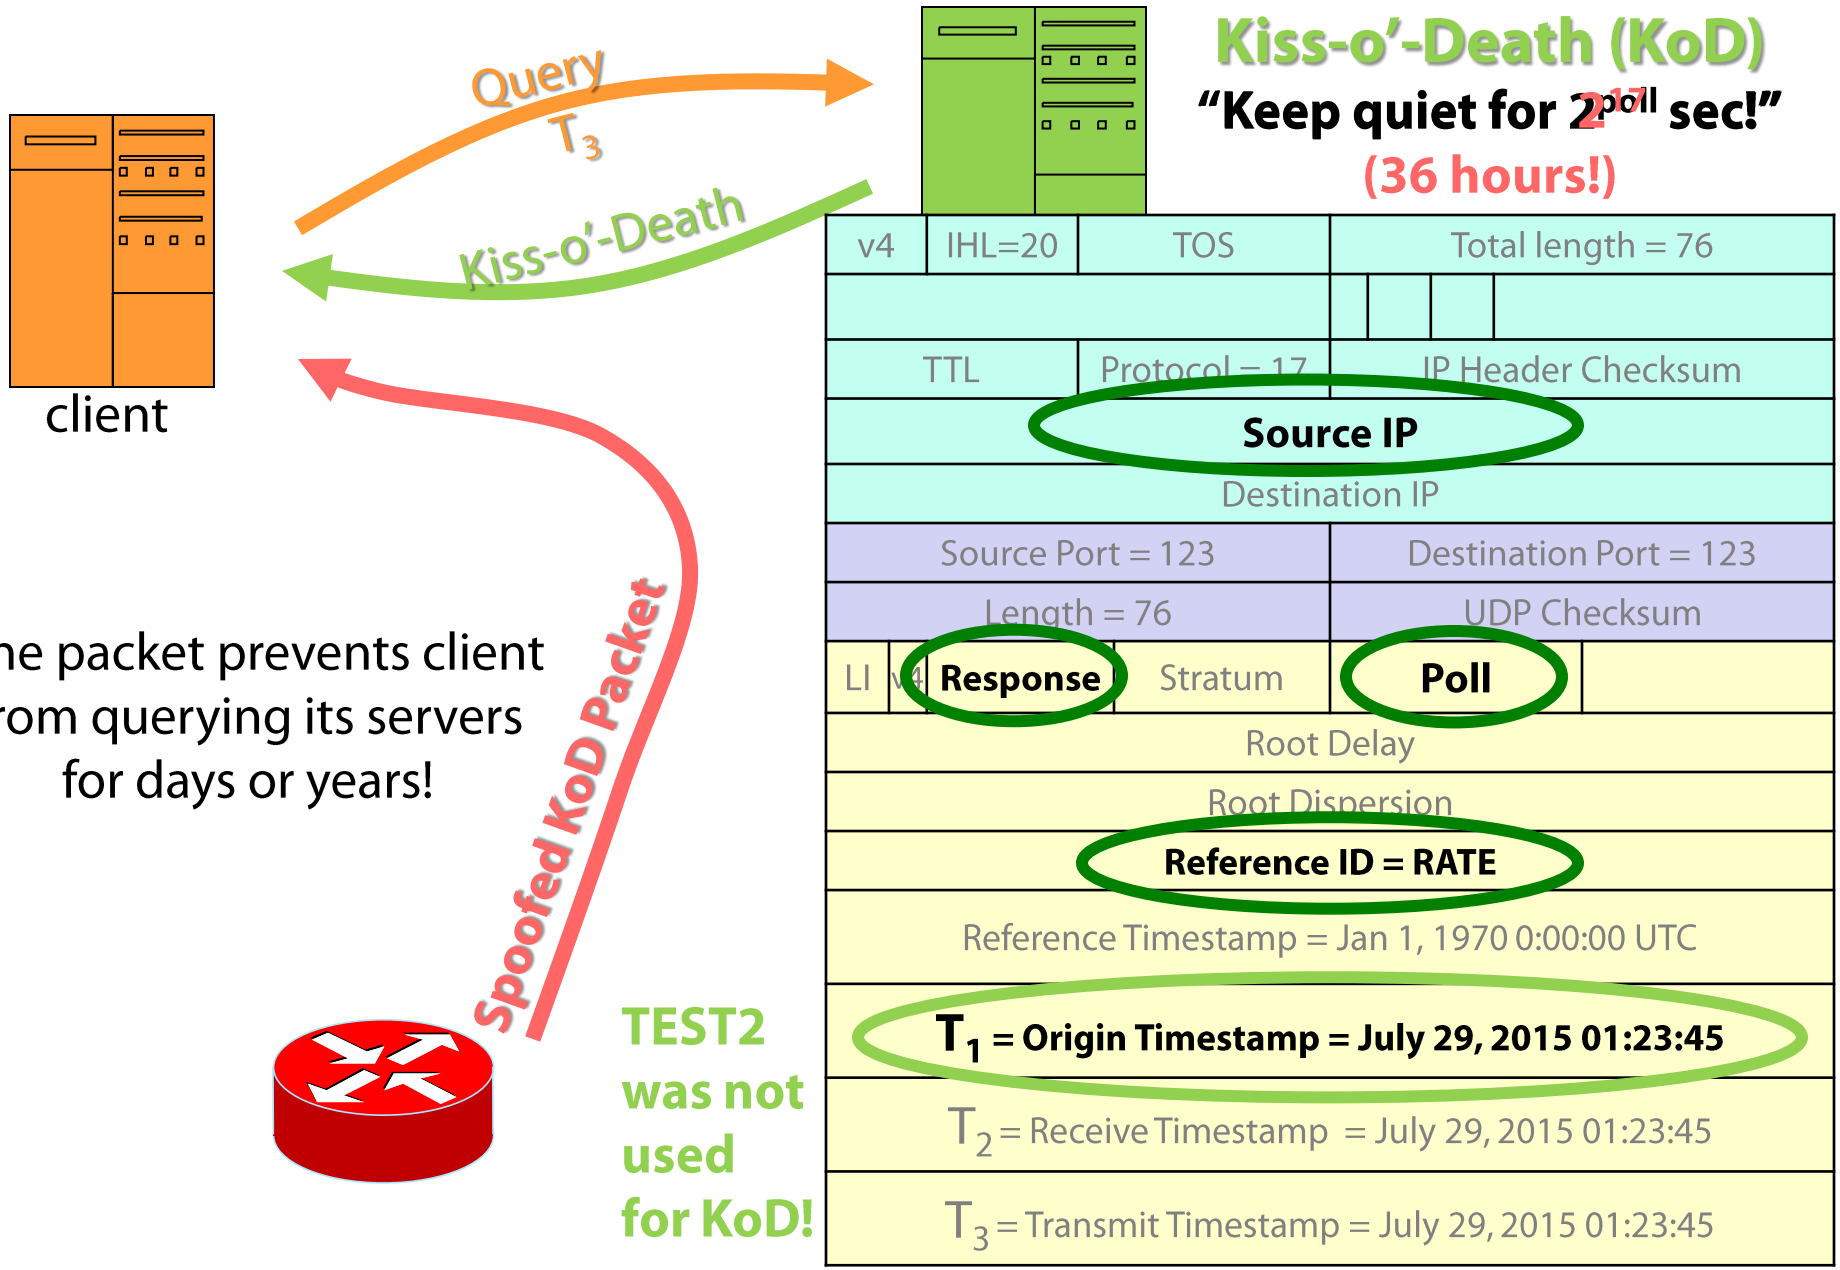

how to learn the server's IP for the spoofed KoD?

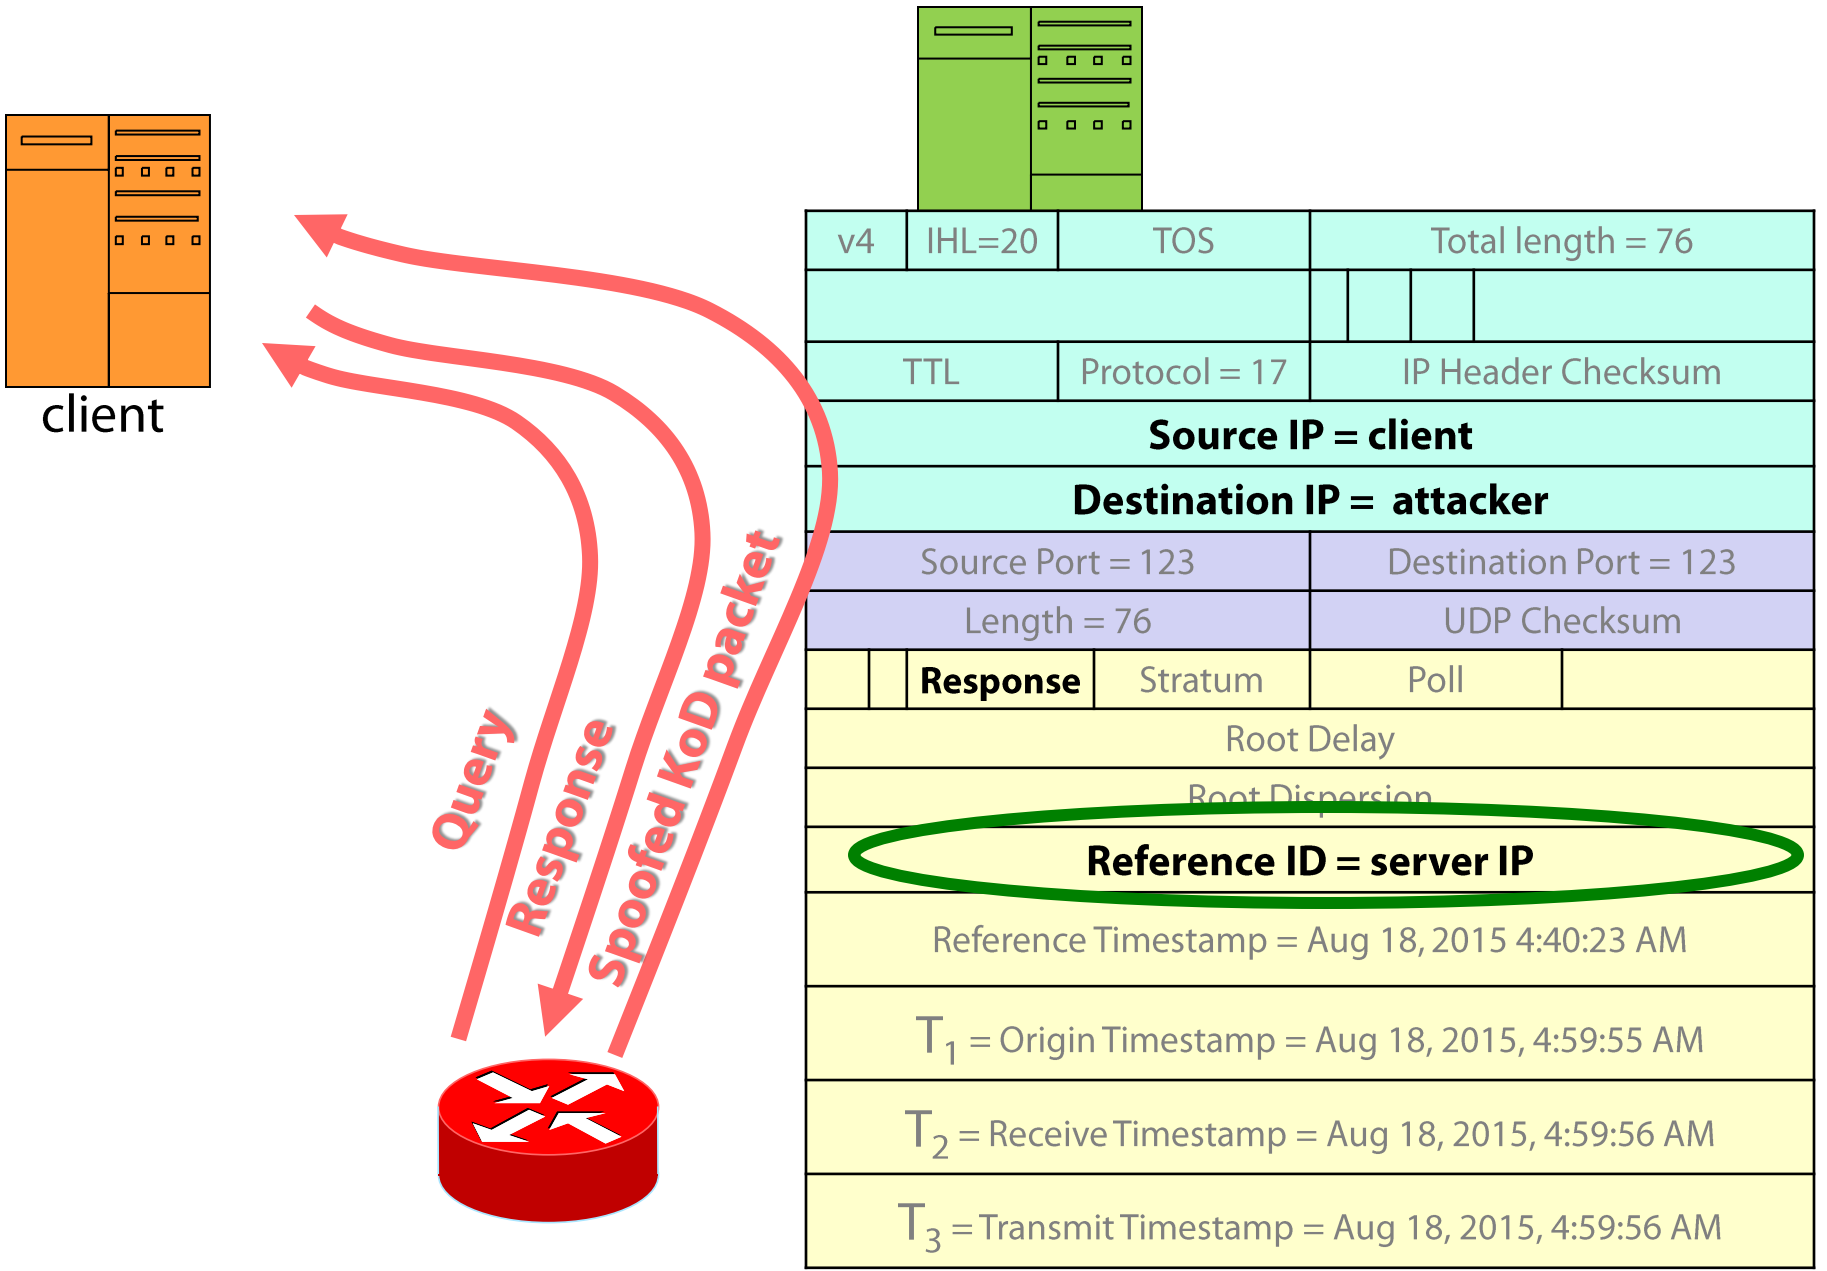

denial of service by priming the pump

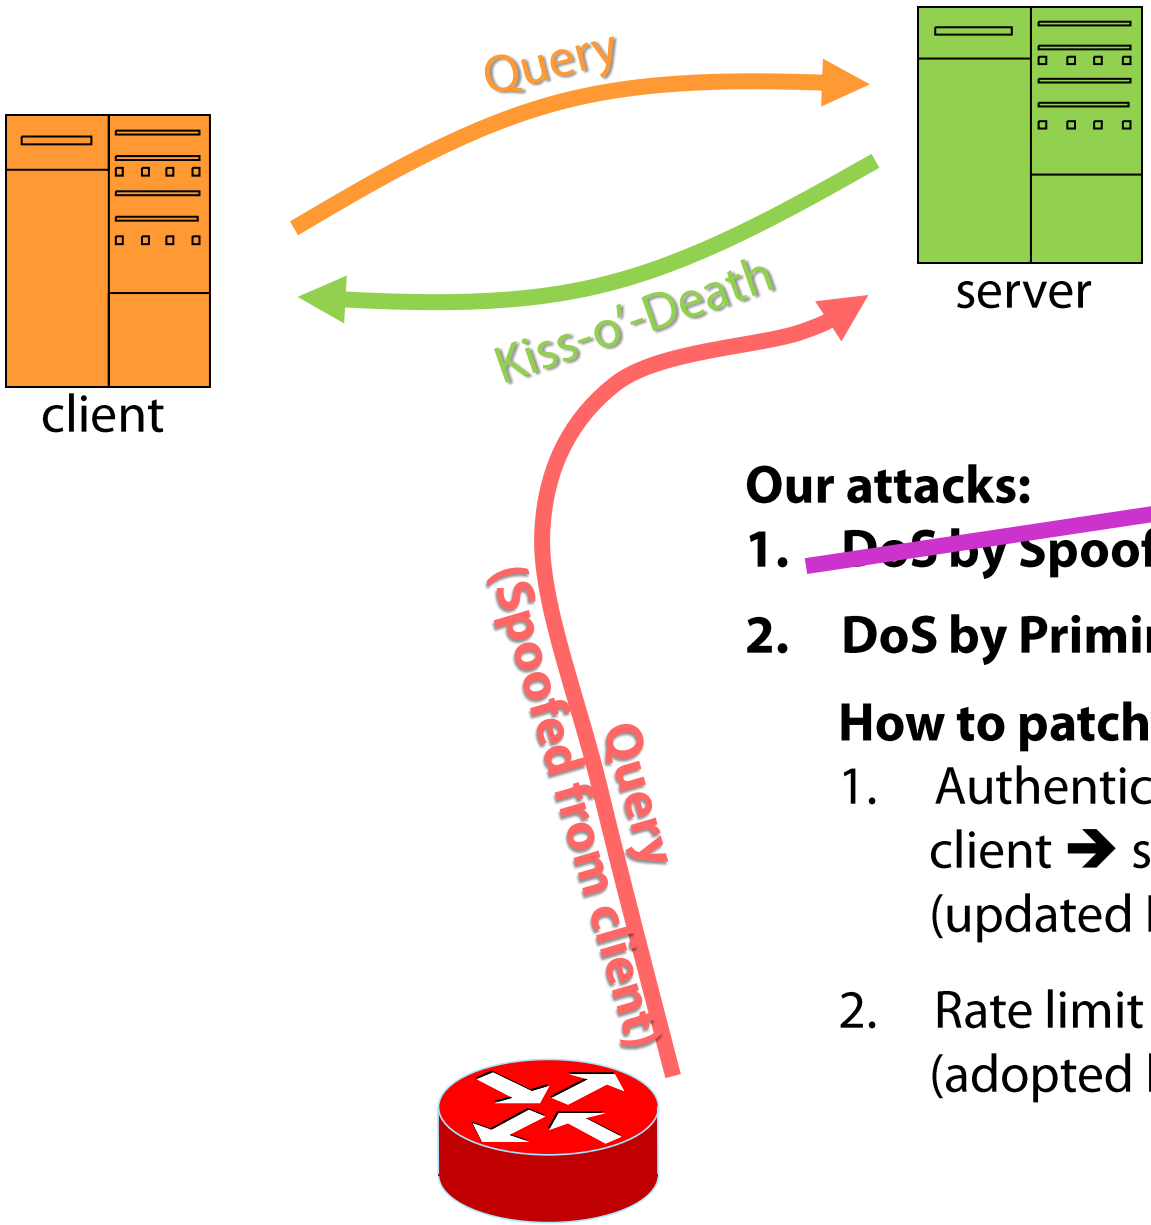

Patched!
TEST2 for KoD
ntpd 4.2.8p4

Our attacks:

1. ~~DoS by Spoofed KoD (off-path)~~
2. **DoS by Priming the Pump (off-path)**

How to patch?

1. Authenticate both directions
   client ➔ server & server ➔client
   (updated Network Time Security IETF ID)
2. Rate limit like DNS Response Rate Limit'g
   (adopted by chrony, ntpd)

outline of the talk

- Background
  - How does NTP work?
  - How does NTP client take time?
- Our client/server mode attacks
  - Denial of Service by Spoofed Kiss-of-Death (off-path)
  - Denial of Service by Priming the Pump (off-path)
  - **Timeshifting** by IPv4 packet fragmentation (off-path)
- Broadcast mode attacks

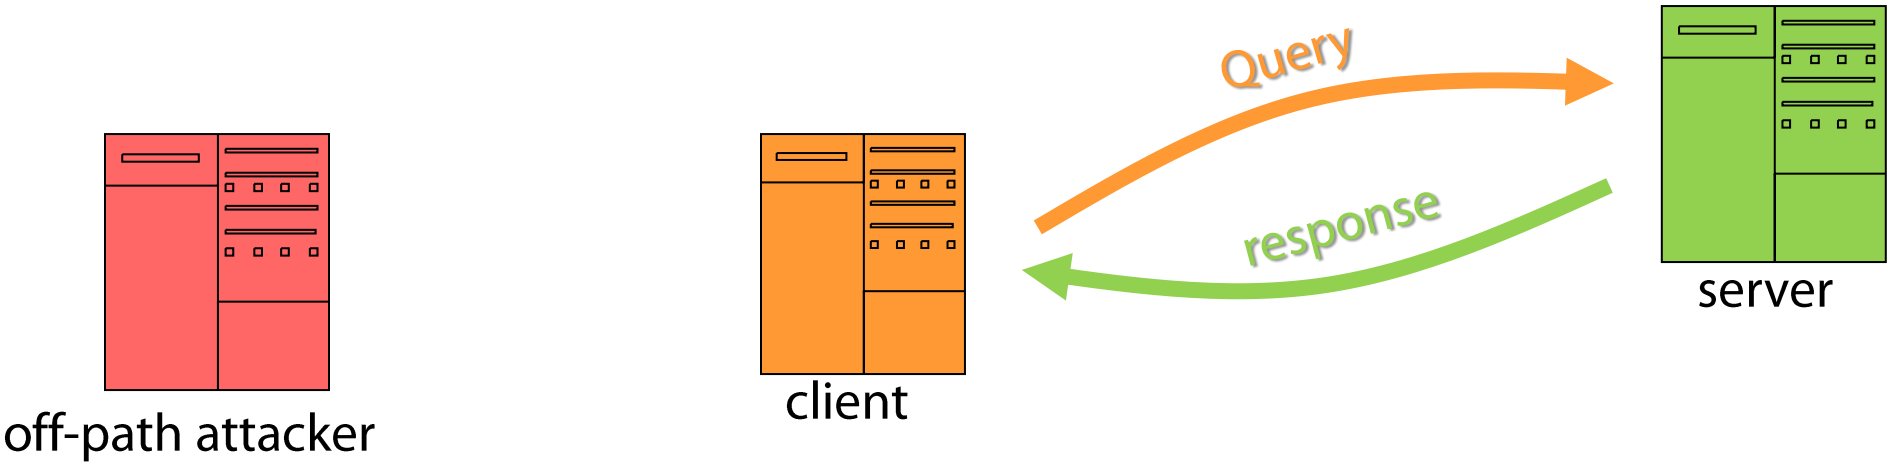

background: IPv4 fragmentation

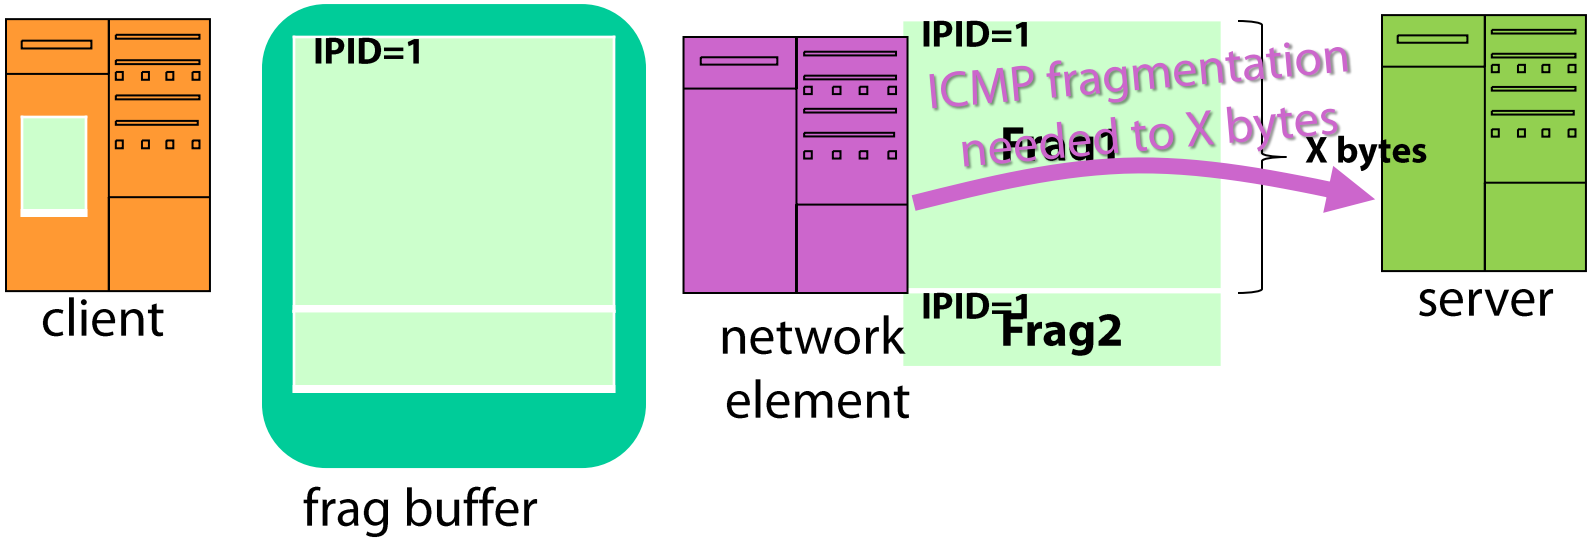

exploiting IPv4 fragmentation to attack NTP

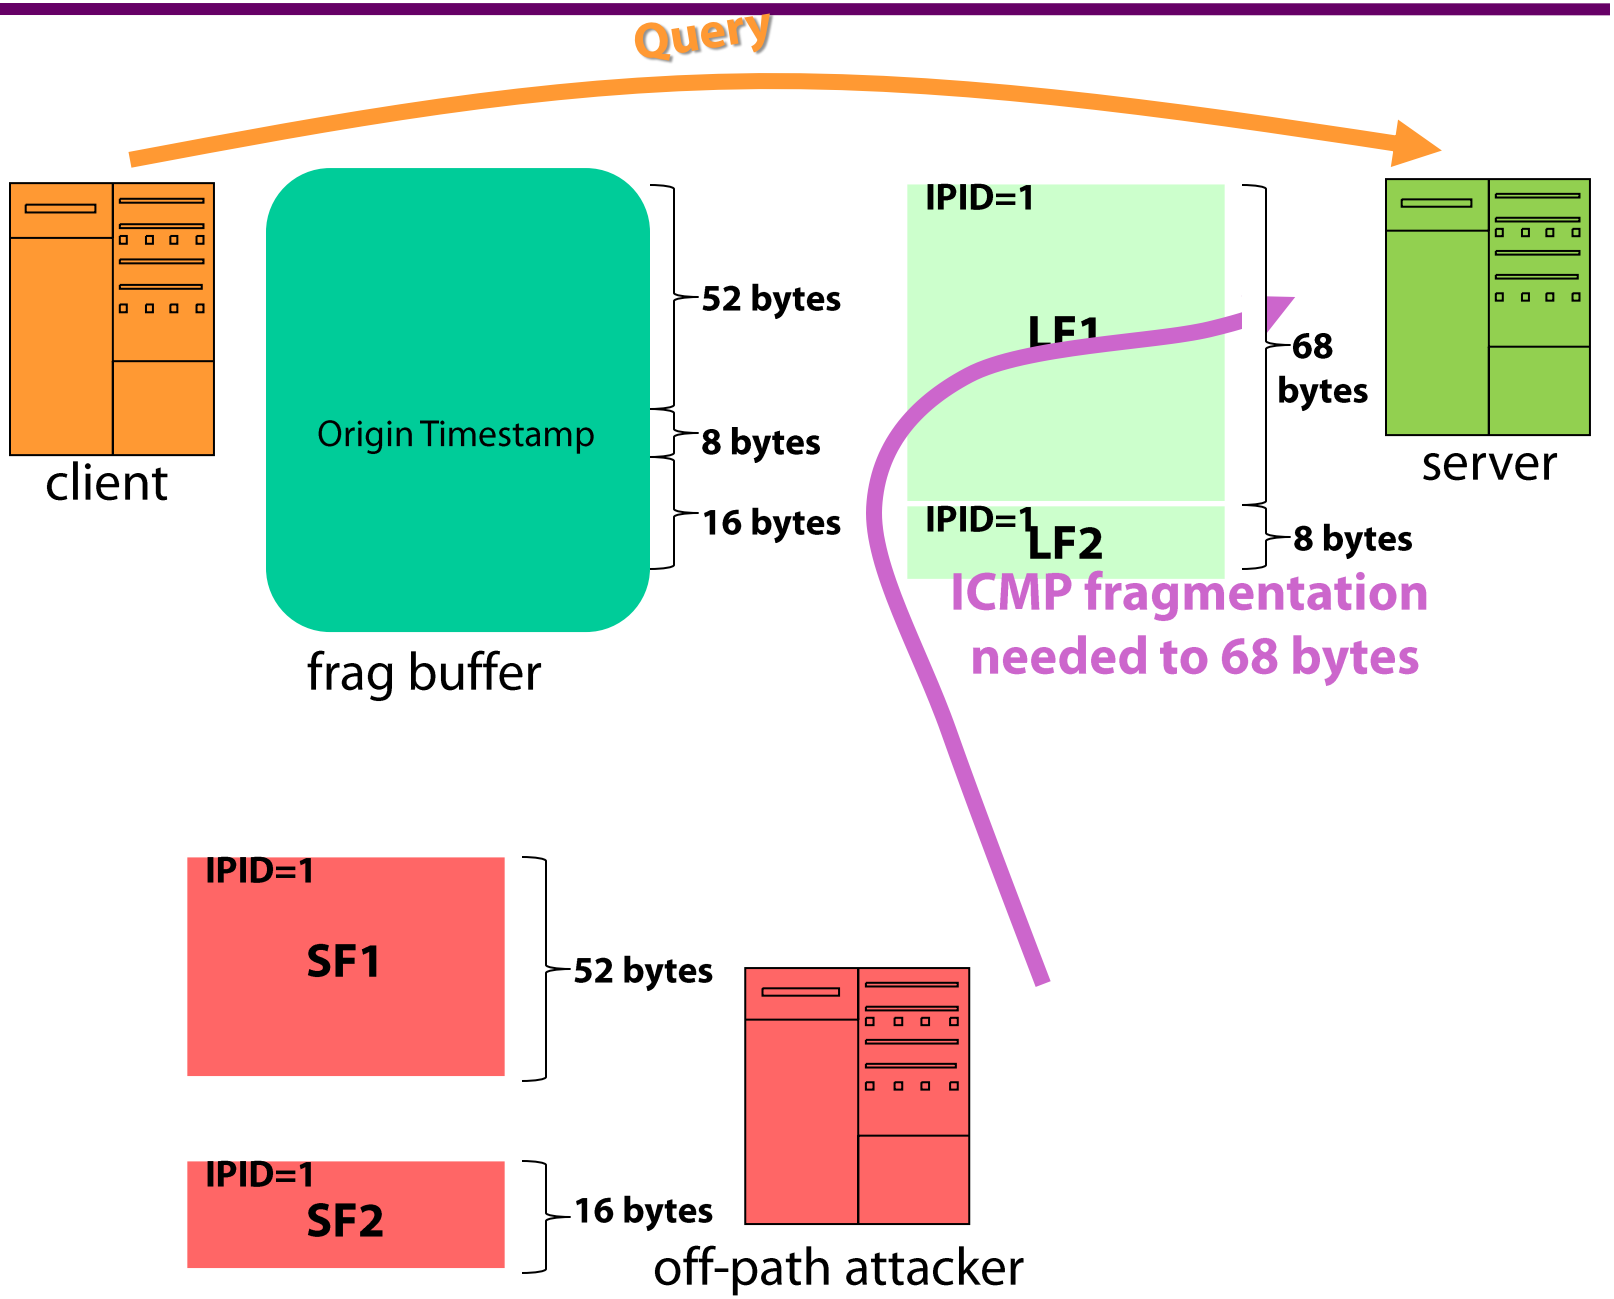

what does the reassembled packet look like?

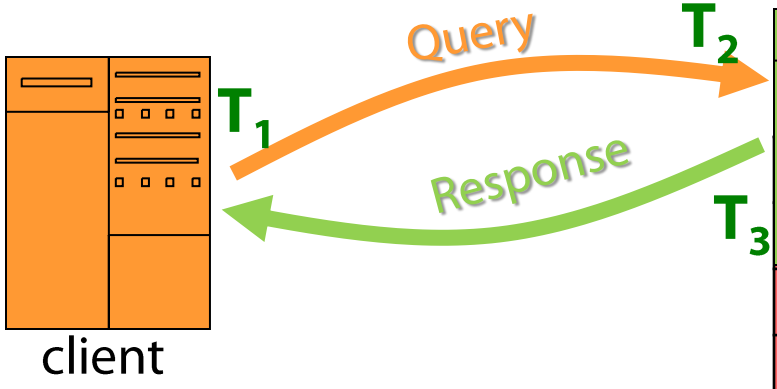

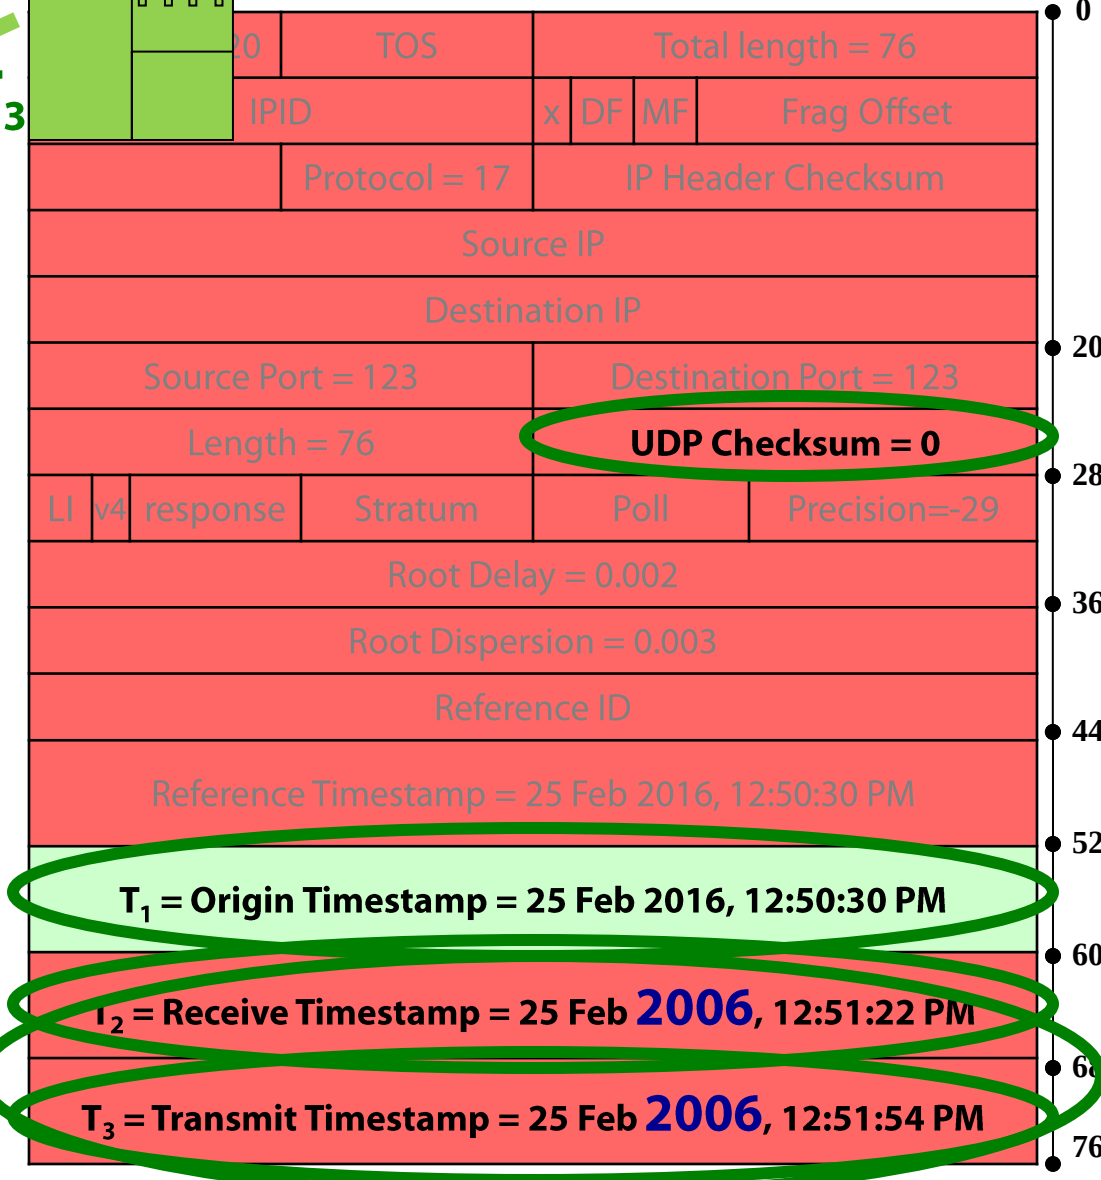

$T_2 - T_1$ = - 10 years + 52 sec

Key Challenge: **Pass TEST2!**

Craft a stream of packets where **T_2-T_1** is consistent within 1 sec!

challenge: construct a stream of consistent packets

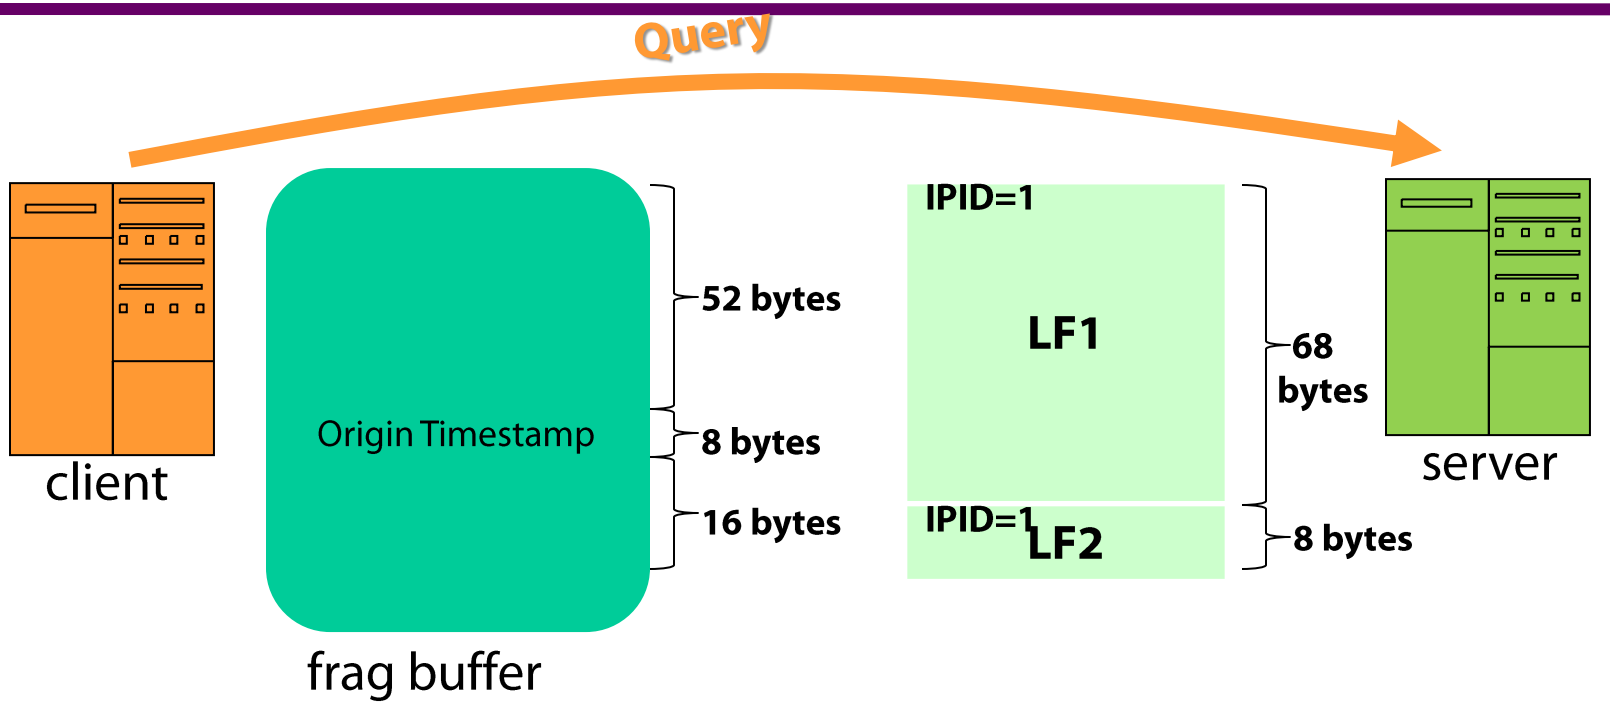

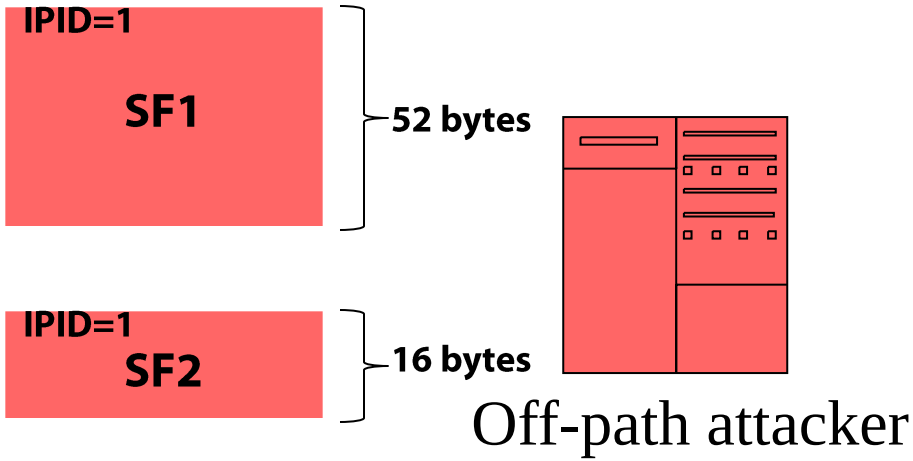

challenge: construct a stream of consistent packets

Why does this help?

The second spoofed fragment sits in the fragment buffer for no longer than 1 sec

$T_2 - T_1 = -10$ years + 1sec

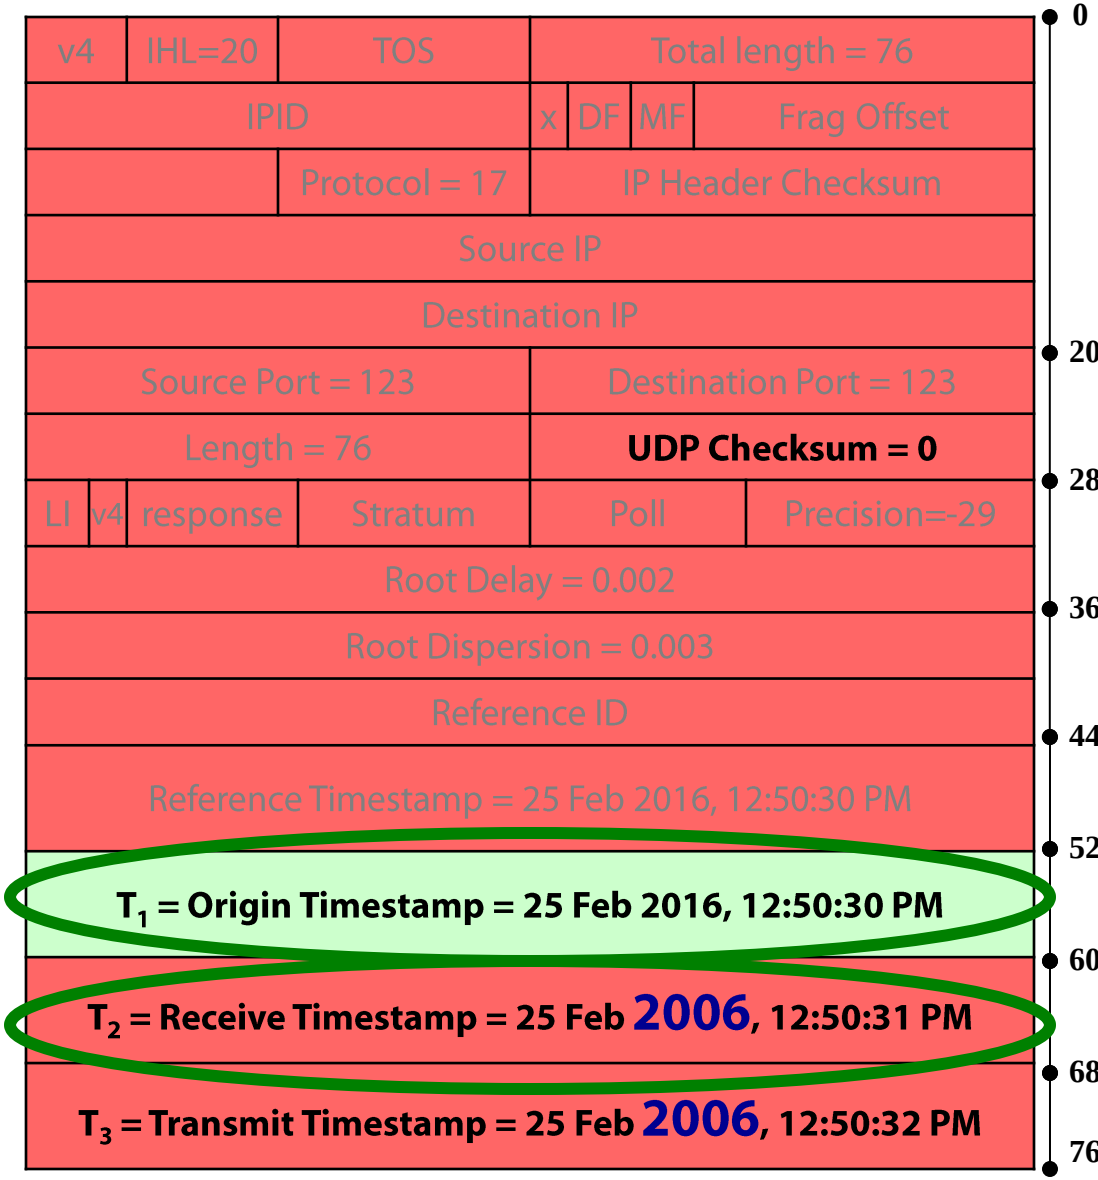

attack surface for our NTP fragmentation attack

Conditions:

1. Server fragments NTP packets to 68 bytes
   - Out of 13M scanned NTP servers, 24K servers do this

2. Client reassembles overlapping fragments by the "First" policy
   - Cannot safely measure due to teardrop [CA-1997-28]

3. Server uses incrementing IPID
   - Inferring globally-incrementing IPID is trivial (most vuln servers)
   - Infer per-destination IPID with **[Gilad-Herzberg'13]** and **[Knockell-Crandall'14]**

Recommendations:

- Servers should not fragment to 68 bytes (Test servers on our site!)
- Drop overlapping IPv4 fragments!

outline of the talk

- **Background**
  - How does NTP work?
  - How does NTP client take time?

- **Our client/server mode attacks**
  - Denial of Service by Spoofed Kiss-of-Death (off-path)
  - Denial of Service by Priming the Pump (off-path)
  - Timeshifting by IPv4 Packet Fragmentation (off-path)

- **Broadcast mode attacks**
- **Other attacks (if time)**

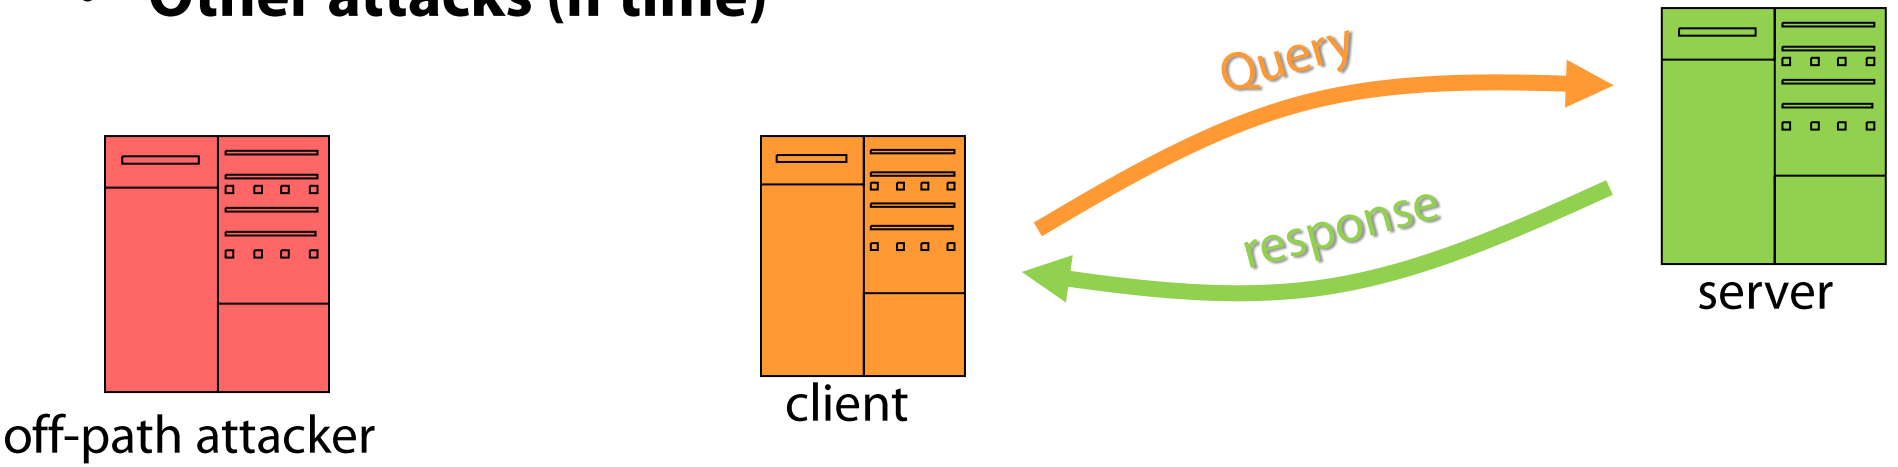

background: broadcast mode

RFC5905 requires broadcast mode to be crypto authenticated

But why?

- client does not apply

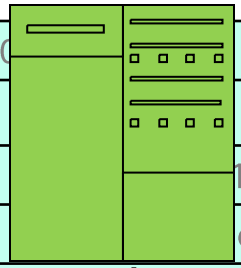

Broadcast server

RFC5905 says:

"If preconfigured to accept broadcast packets, the client accepts packets from ANY server that sends it broadcast packets."

how does a broadcast client detect replay attacks?

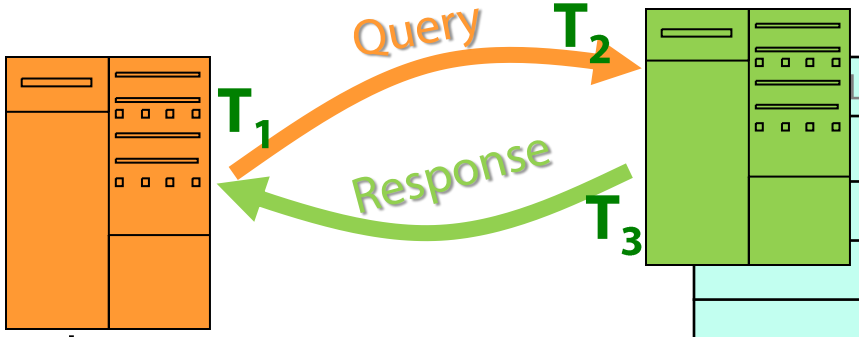

L=20	TOS	Total length	IP header
	Protocol = 17	IP Header Checksum	
Source IP			
Destination IP			
Source Port = 123		Destination Port = 123	UDP header
Length = 76		UDP Checksum	
LI v4	Broadcast Stratum	Poll	
Root Delay			NTP data
Root Dispersion			
Reference ID			
Reference Timestamp = Jan 1, 1970 0:00:00 UTC			
T_1 = Origin Timestamp = NULL			
T_2 = Receive Timestamp = NULL			
T_3 = Transmit Timestamp = March 22, 2016 01:53:45			
Key ID= 00000001			NTP MAC
Message Digest = 324a4b23130fff3eab4581931ee6fa5d4			

TEST1: Prevents replay of **most recent** packet!

déjà vu: time sticking attack via packet replay

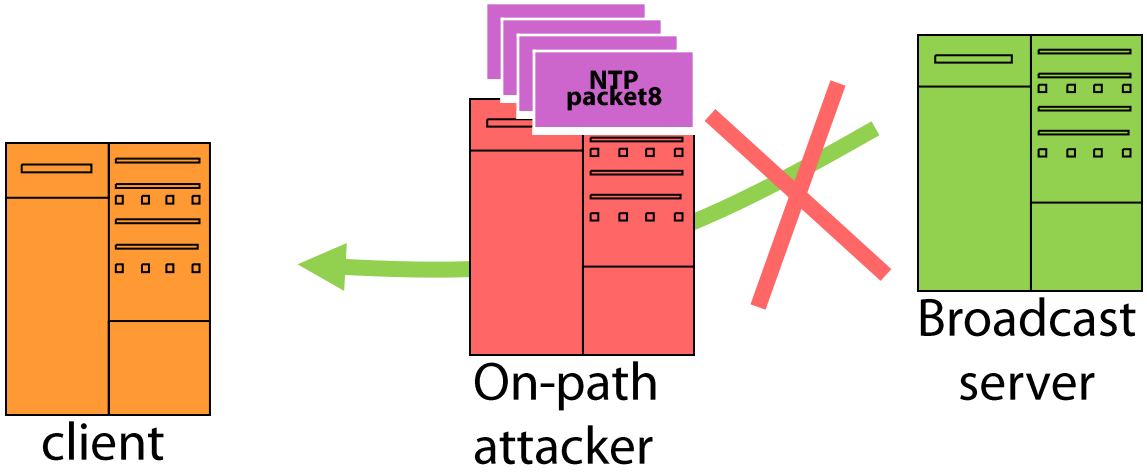

Attack: Replay, in order the eight most recent timestamps

Recommendation:

Add an incrementing counter to the (timestamp) fields that are null

off-path DoS attack using malformed crypto packets

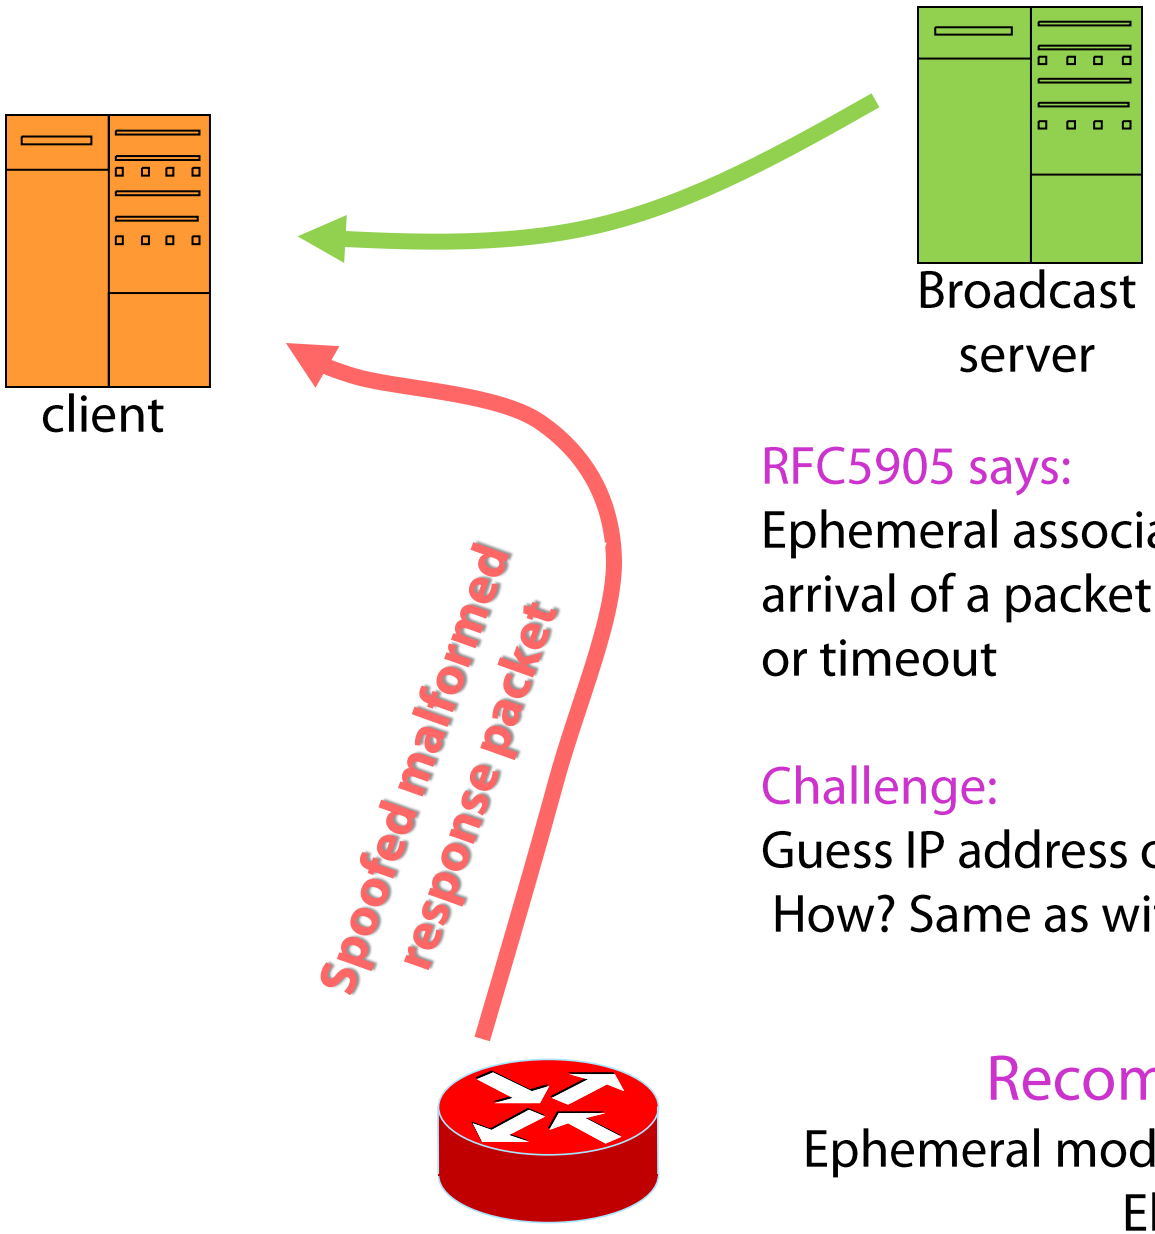

RFC5905 says:

Ephemeral associations are mobilized upon arrival of a packet and demobilized upon error or timeout

Challenge:

Guess IP address of the server.

How? Same as with our KoD attacks

Recommendation:

Ephemeral modes considered harmful. Eliminate.

other recent attacks

The dreaded NTP DDoS Reflection attack

- this is still a problem
- send the monlist control query to an NTP host (via UDP!)
- … get a list of last 600 IPs interacting with that host.

Cisco ASIG attacks on NTP

- zero origin timestamp attack (from RFC5905)
- NAK to the future (crypto NAK implementation flaw)
- crypto key misbinding leading to a sybil attack vulnerability
- origin timestamp leak vulnerability
  - Use NTP control queries to learn exact value of origin timestamp
  - … and bypass TEST2.

recommendations for NTP users

- **Firewall your NTP instances**.
  - End-hosts should not accept **anything** other than mode 4 packets from their preconfigured servers.
  - All hosts (incl. servers) should not accept **any** control queries (not just monlist) from arbitrary IPs. (Use the ntpd **`noquery`** option)
  - Firewalls should block KoD (mode 4) with high poll (eg poll >10)

- If ntpd is configured with the –g option, monitor for reboots
  - NTP is much more vulnerable to attacks when it reboots

- Don't use broadcast mode except in a safe firewalled network
  - Even if your packets are authenticated, you are still vulnerable.
  - Make sure no broadcast packets come into your network from outside

- Don't fragment NTP packets

longer recommendations for securing NTP

- Stop leaking so much information.
  - NTP packet leak the reference ID and reference time.
  - NTP control packets leak timestamps and lots of internal state
  - All this information can be collected using UDP
  - And in most cases is available by default
- Get rid of the KoD, use RRL style rate limiting instead
- Use a modern control protocol instead of leaky UDP control protocol
- Latest burst of bugs has shown that RFC5905 is underspecified
- Lots more work to be done to develop crypto for NTP

Questions?

Attacking the Network Time Protocol
Aanchal Malhotra, Isaac E. Cohen, Erik Brakke and Sharon Goldberg
NDSS'16, San Diego, CA. Feb 2016.

Attacking the NTP's Authenticated Broadcast Mode
Aanchal Malhotra and Sharon Goldberg
SIGCOMM Computer Communication Review, April 2016.

http://www.cs.bu.edu/~goldbe/NTPattack.html