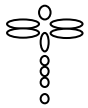

root attack
~ end-user view ~

Matsuzaki ‘maz’ Yoshinobu
<maz@iij.ad.jp>

attack on 6th Feb 2007

- DDoS against the root and some TLDs
  - 10:00UTC～
- Attack Traffic
  - UDP/53
  - large packet
  - Asia is a major source of the attack traffic

reports

- Several reports are published
  - how much the attack traffic...
  - how ops work together...
- Was there any effect for end-users?
  - delay
  - any failure on name resolve

dns cache server in IIJ/AS2497

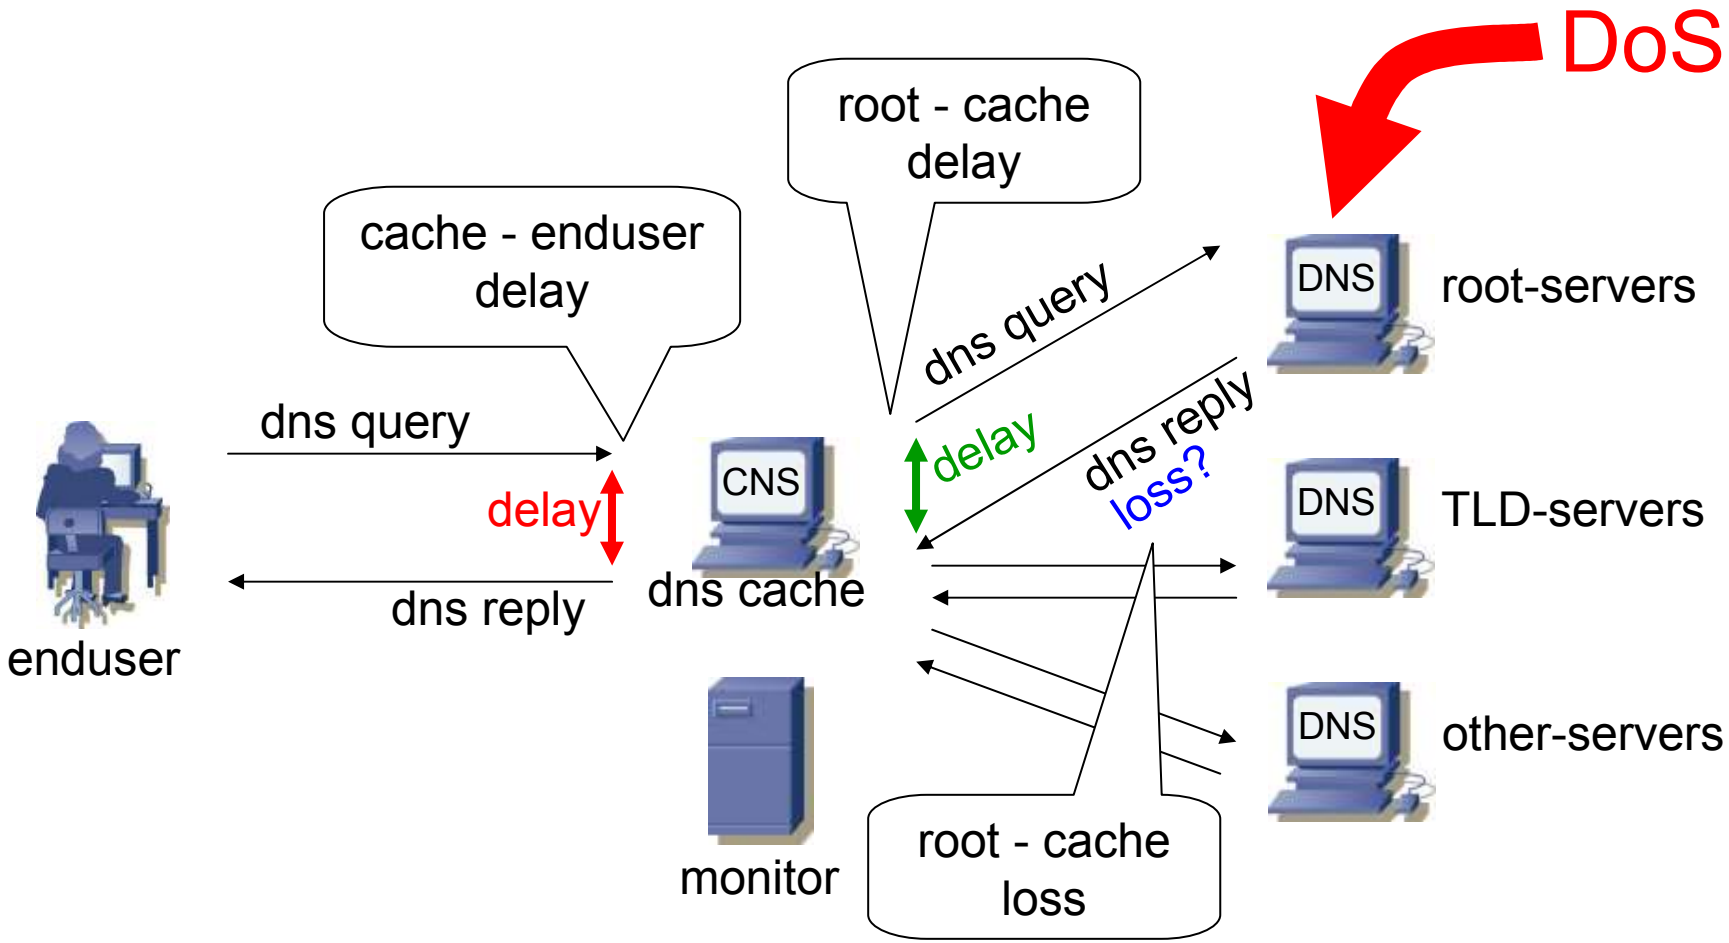

response delay of cache server

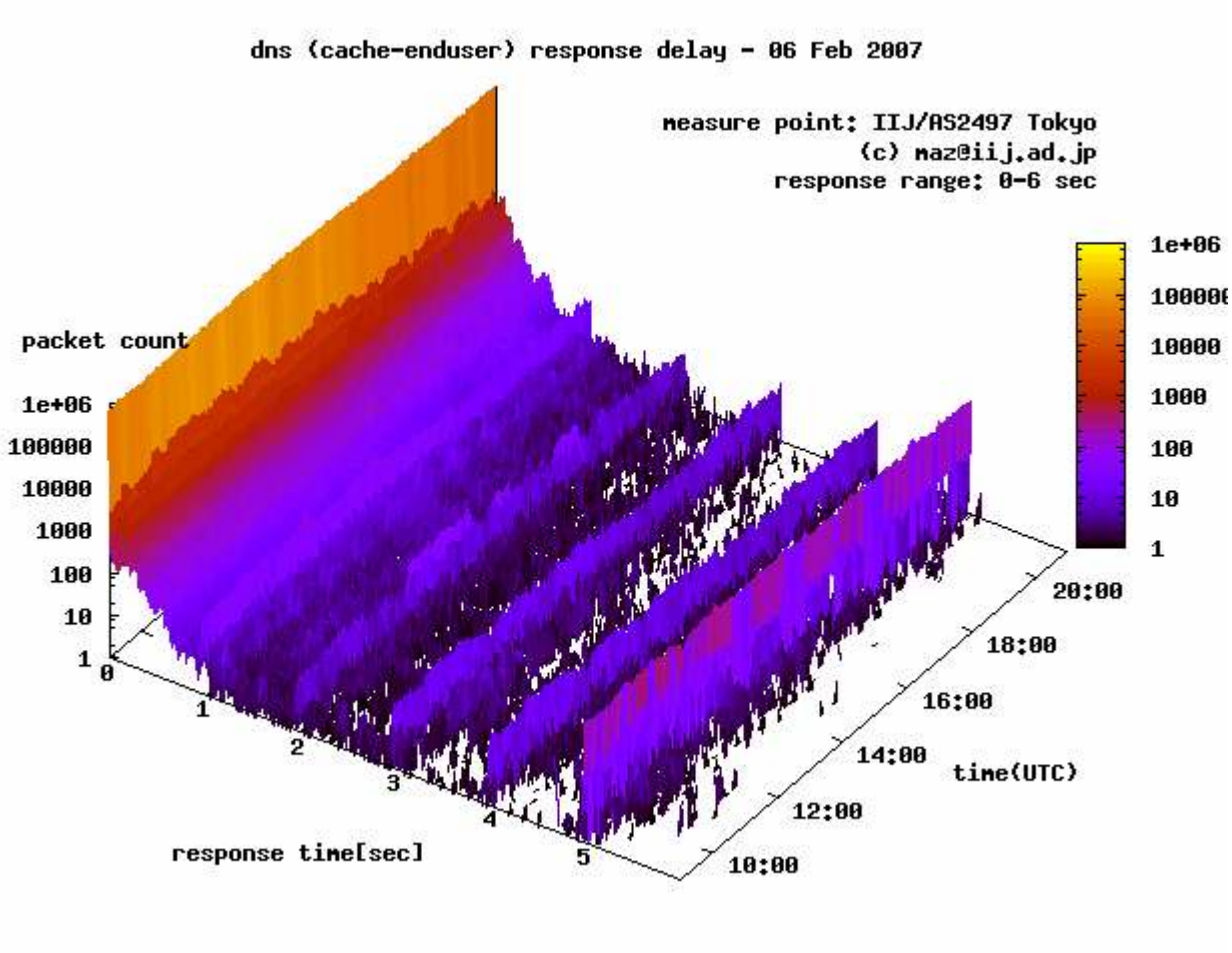

zoom (0~1sec delay)

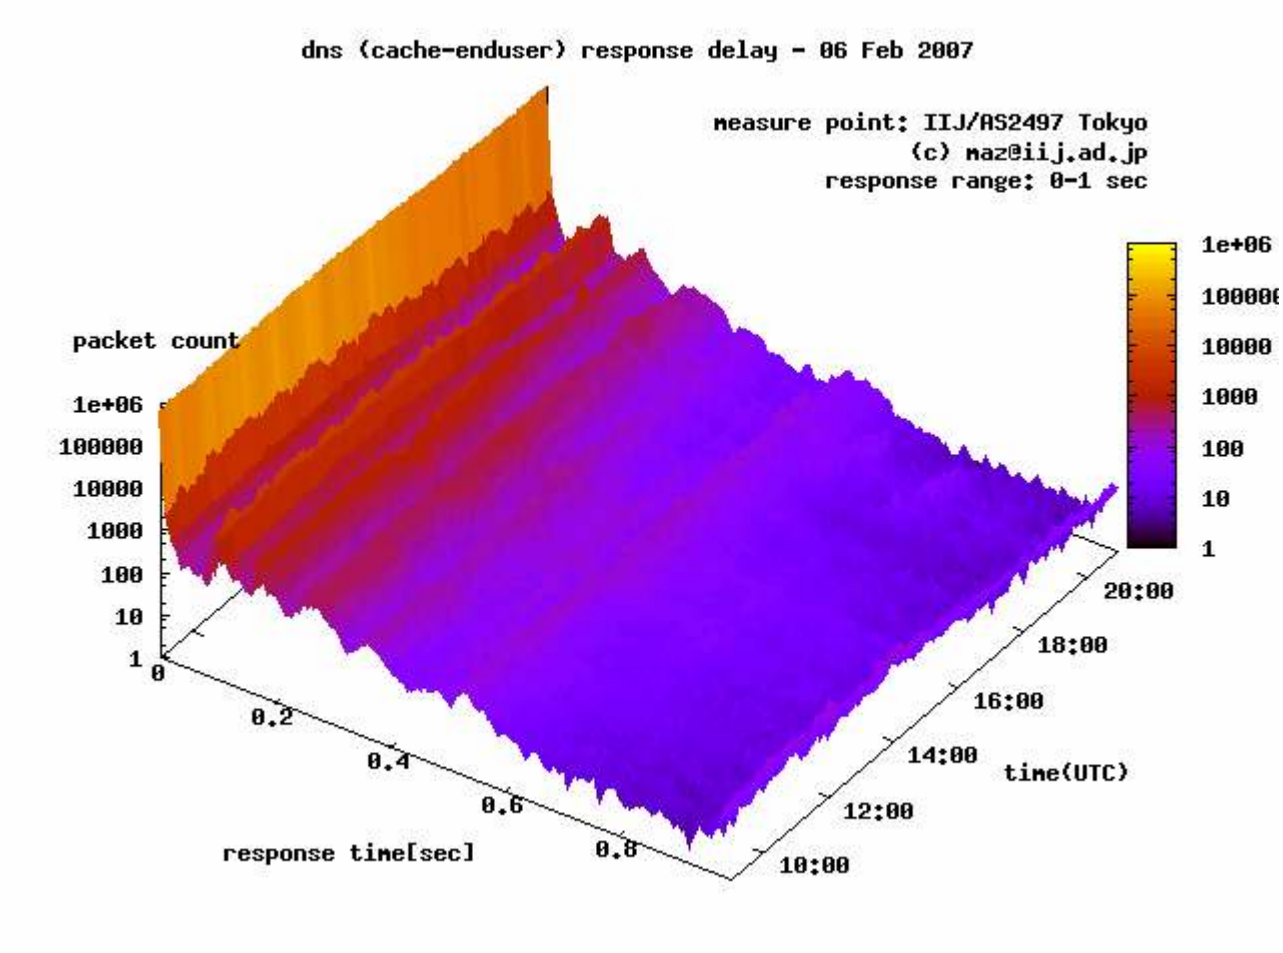

AforA

e.g.) A? 192.168.0.1.

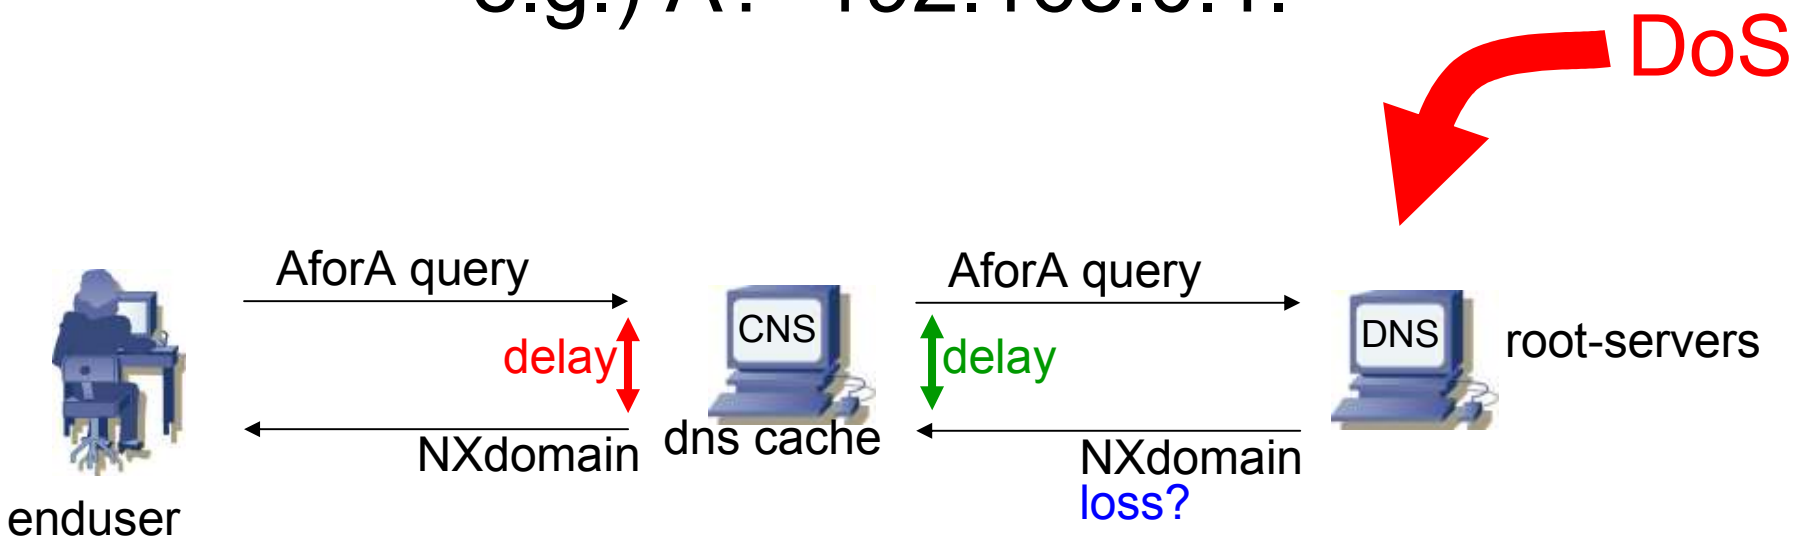

- We can estimate the root-server performance by checking the delay of AforA queries.

AforA query stat

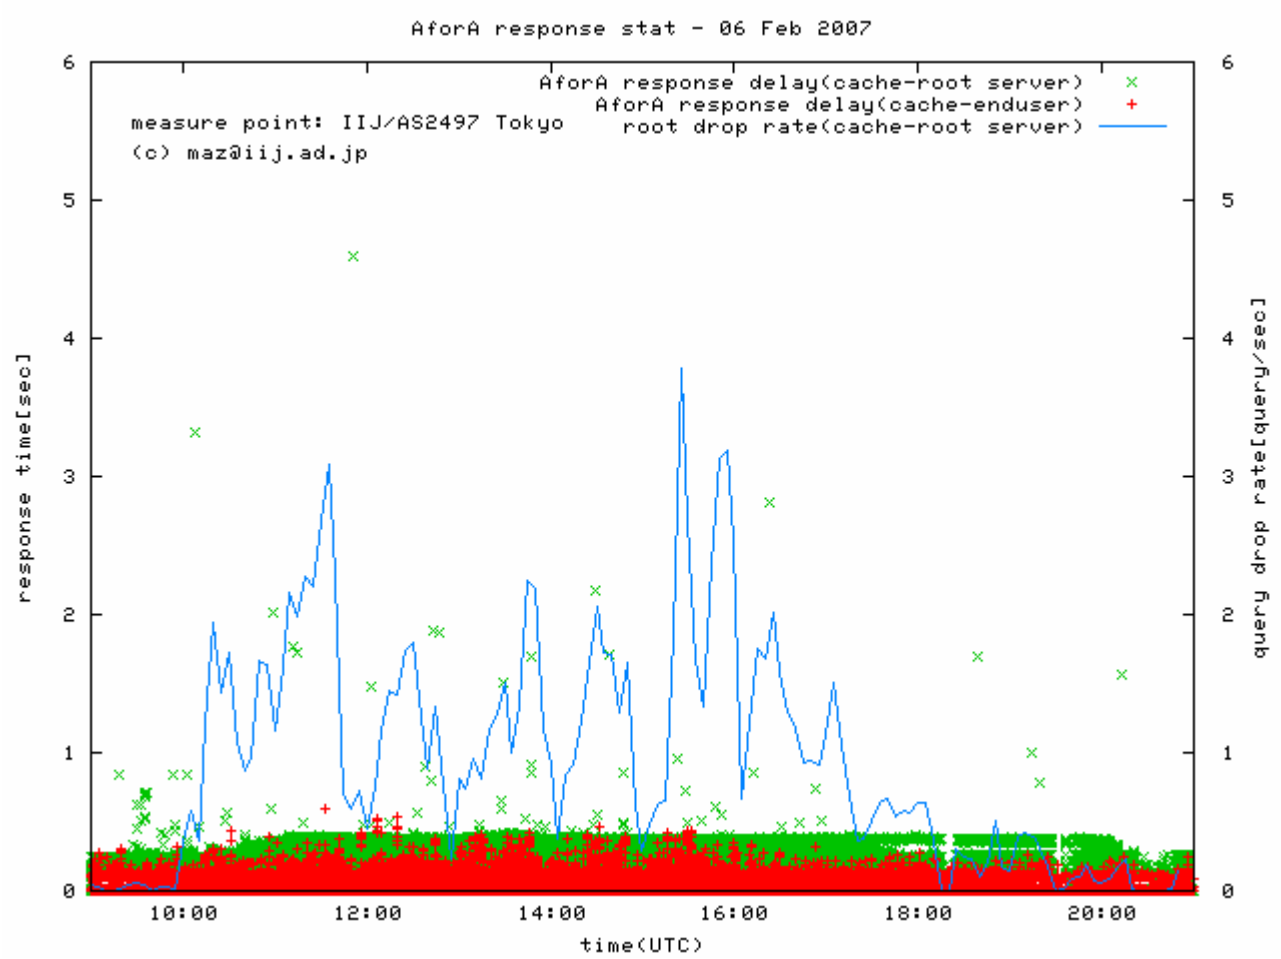

response delay - all recursive query

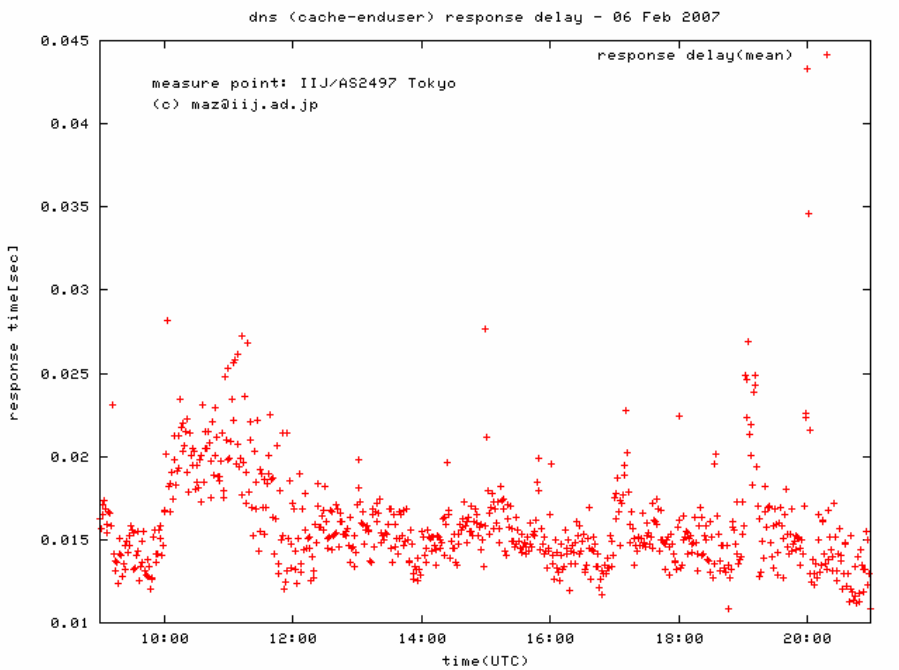

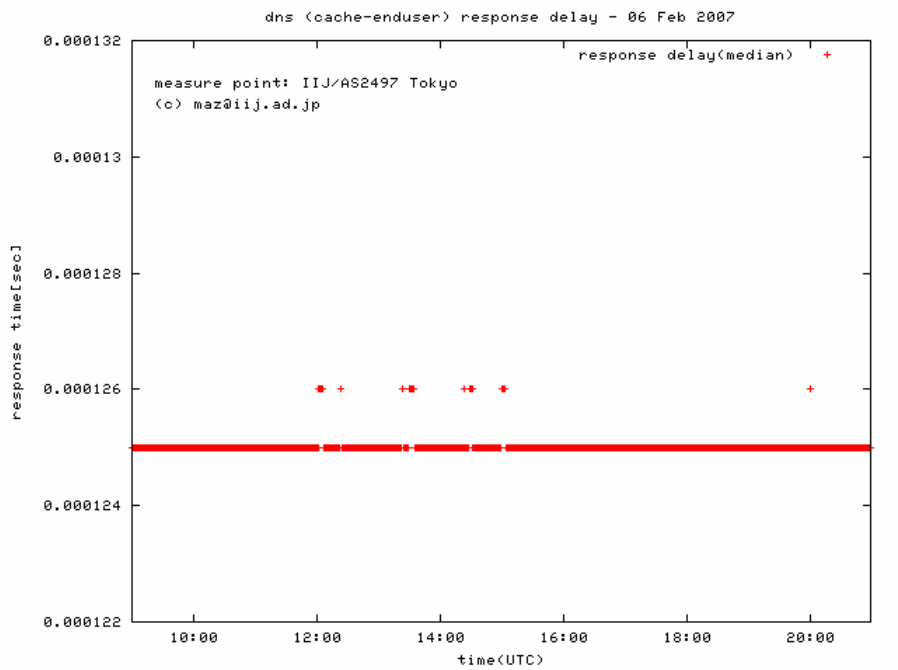

response delay - .uk query only

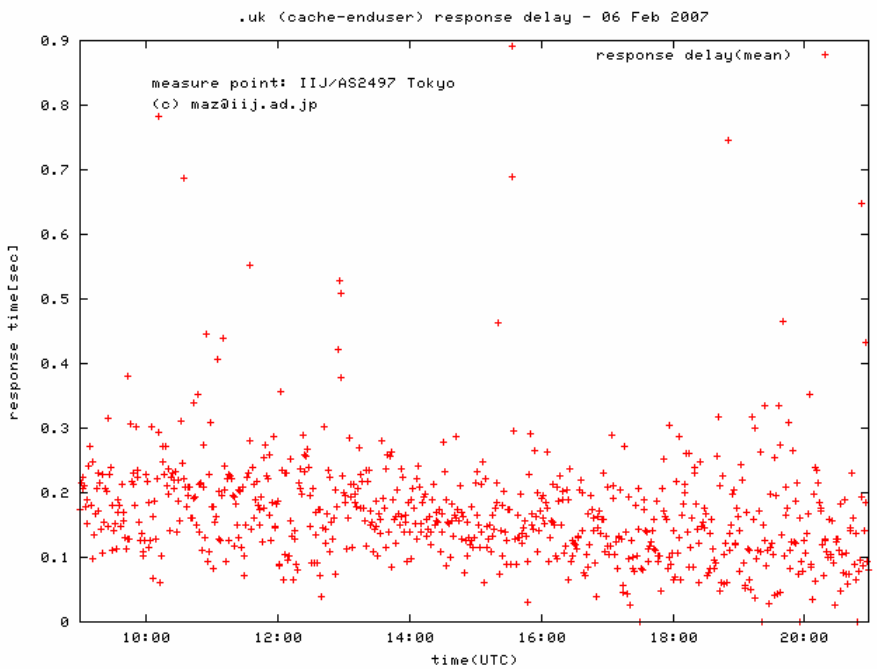

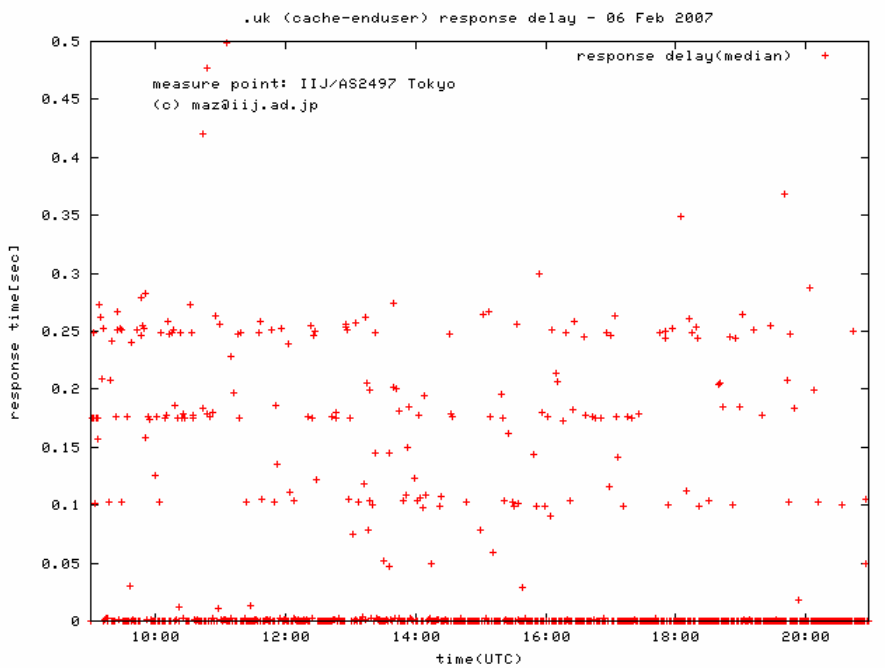

response delay - .org query only

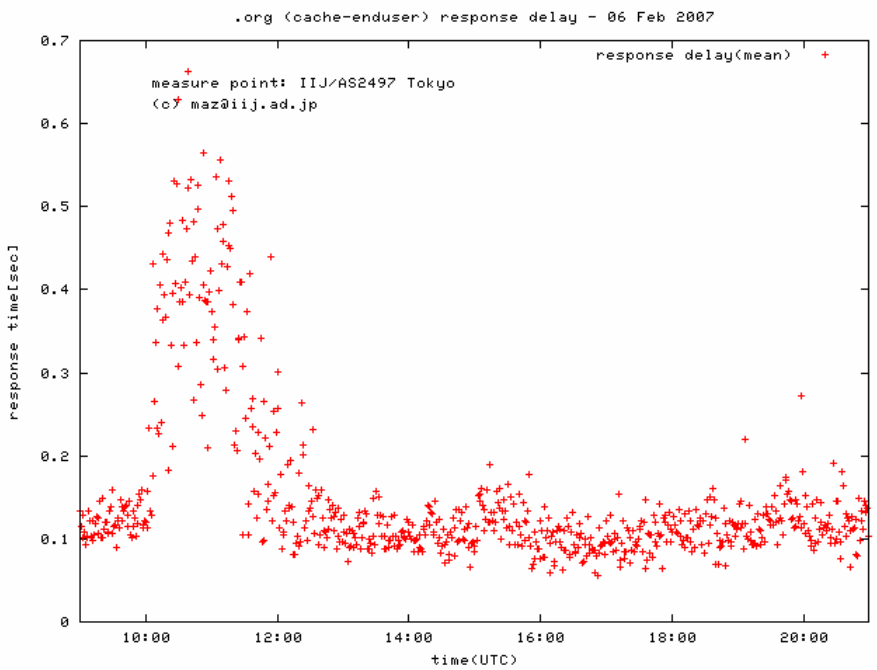

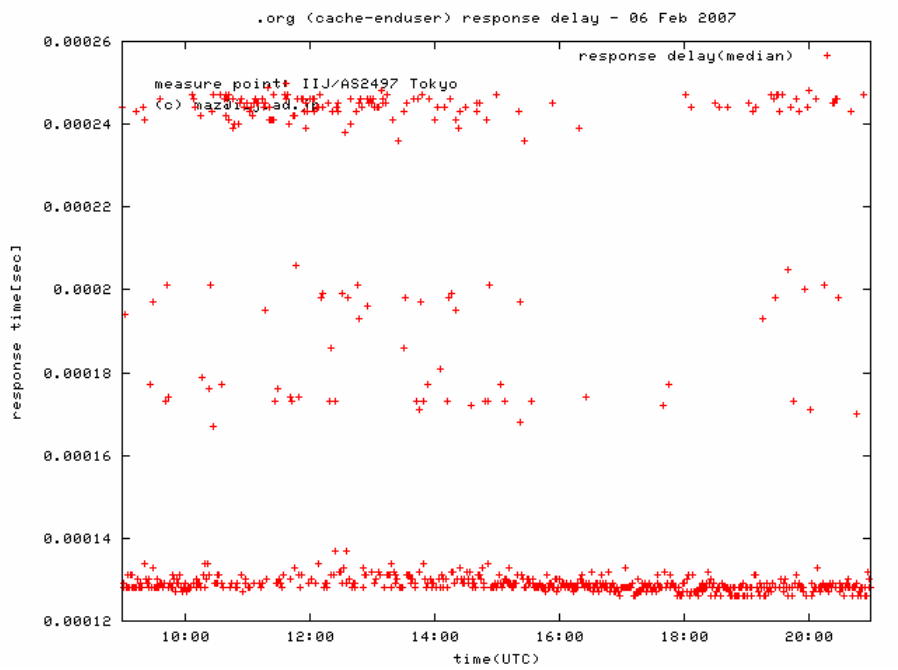

.uk query stat

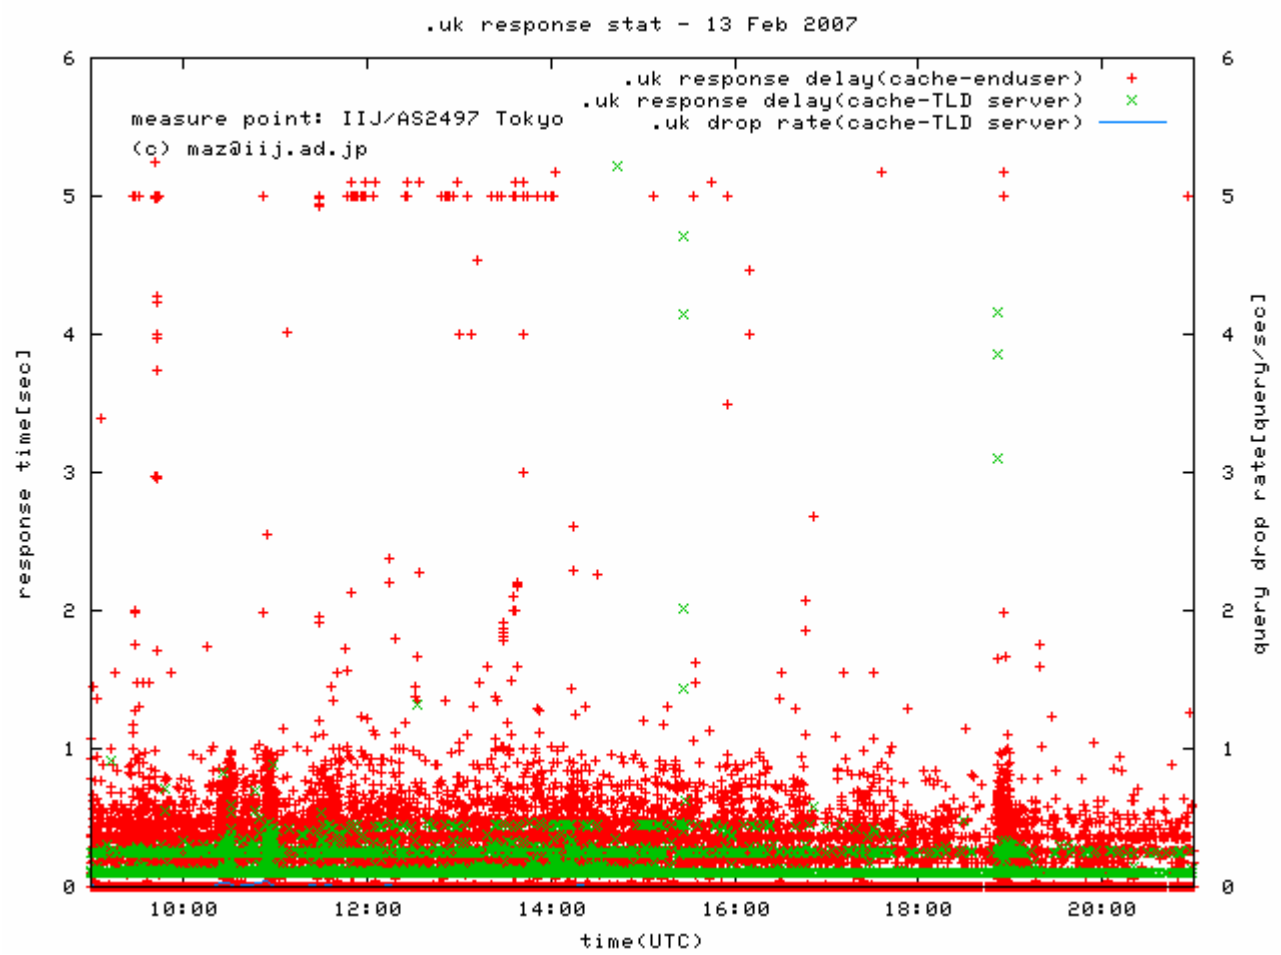

.org query stat

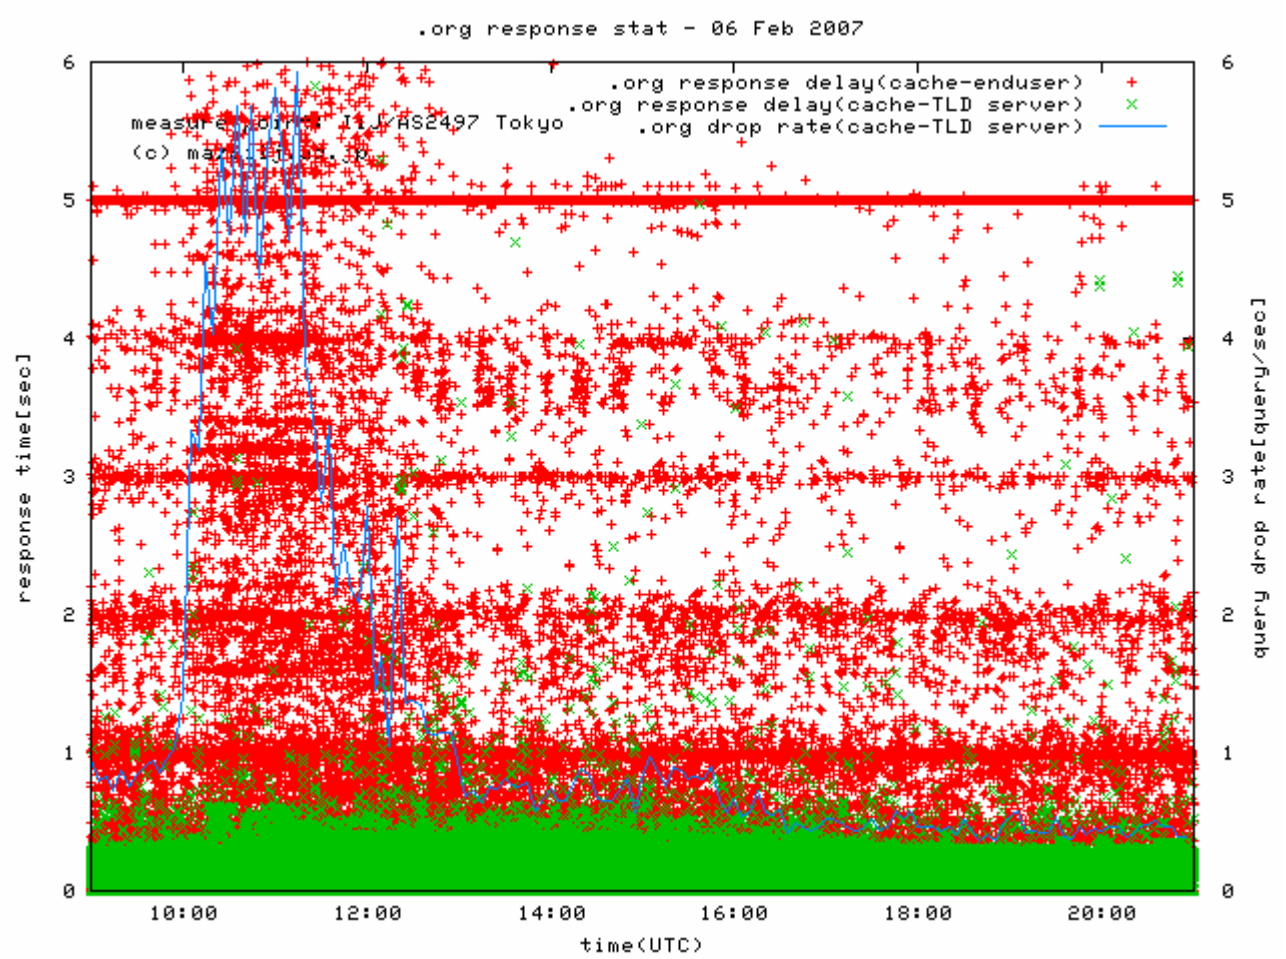

query to root-servers

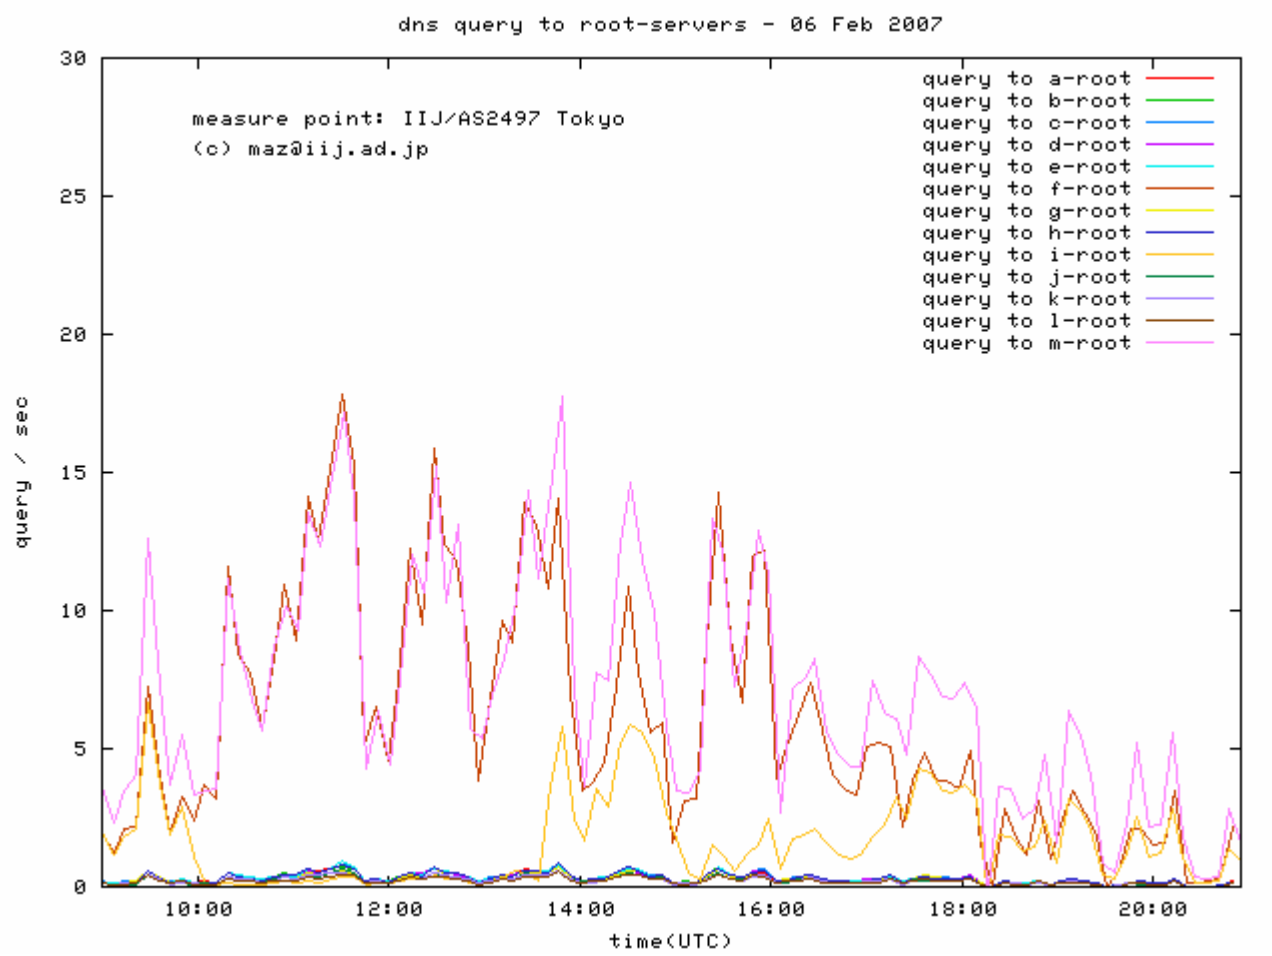

response delay of m.root

hostname.bind. - "M-NRT-JPNAP-3"

during attack

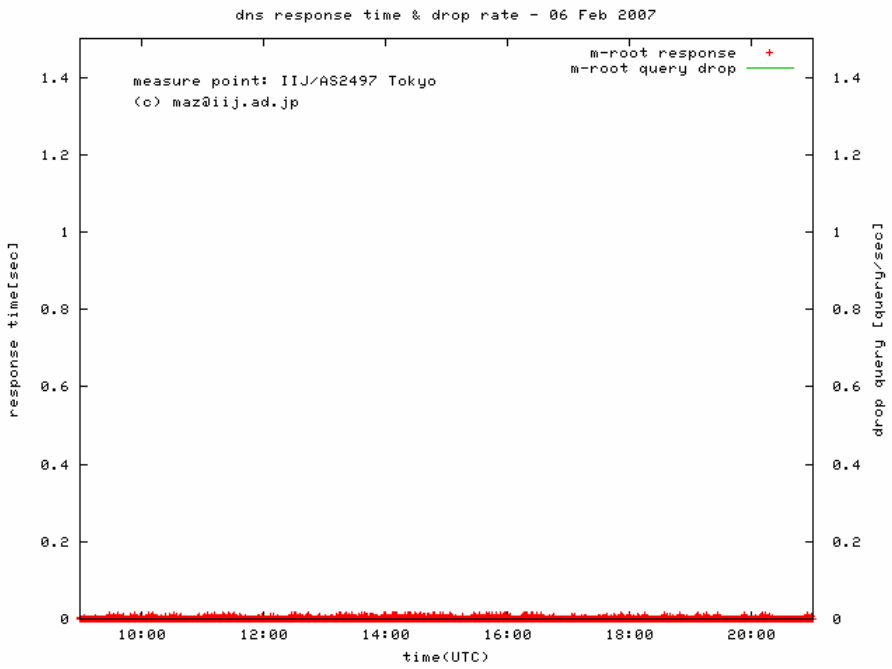

1 week later

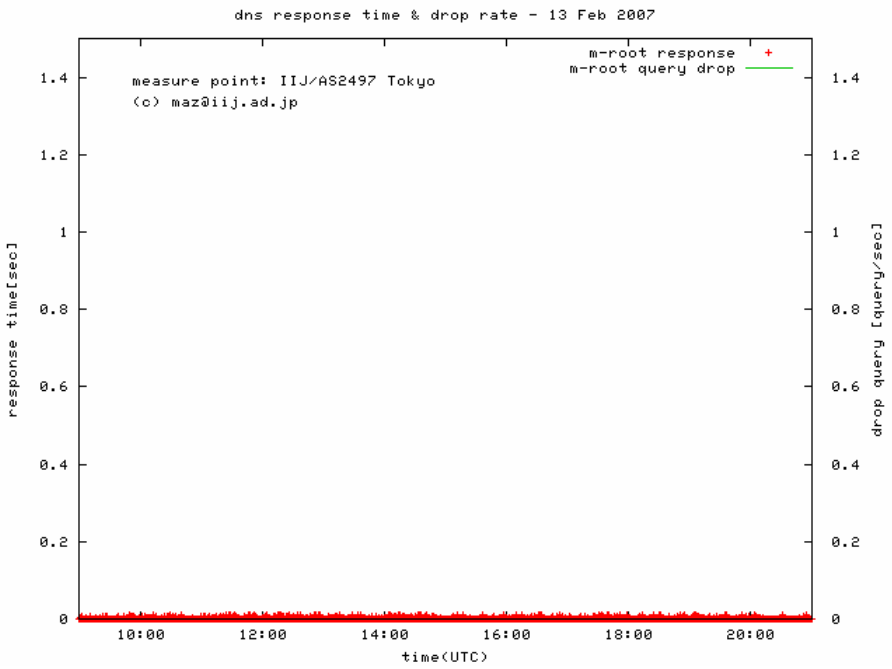

IIJ/AS2497 and m.root

- IIJ have 3 peers with m.root. anycast sites.
  - IIJ provides transit for m.root.

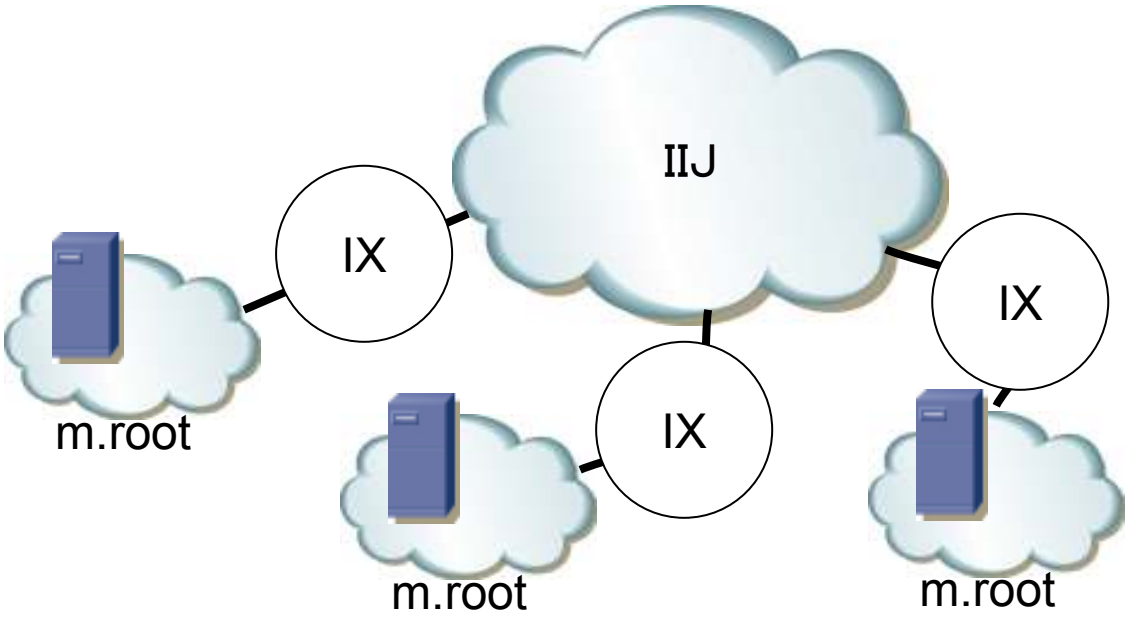

during the attack

- IIJ transited attack traffic as well...
  - IIJ's cache server selected the other site.

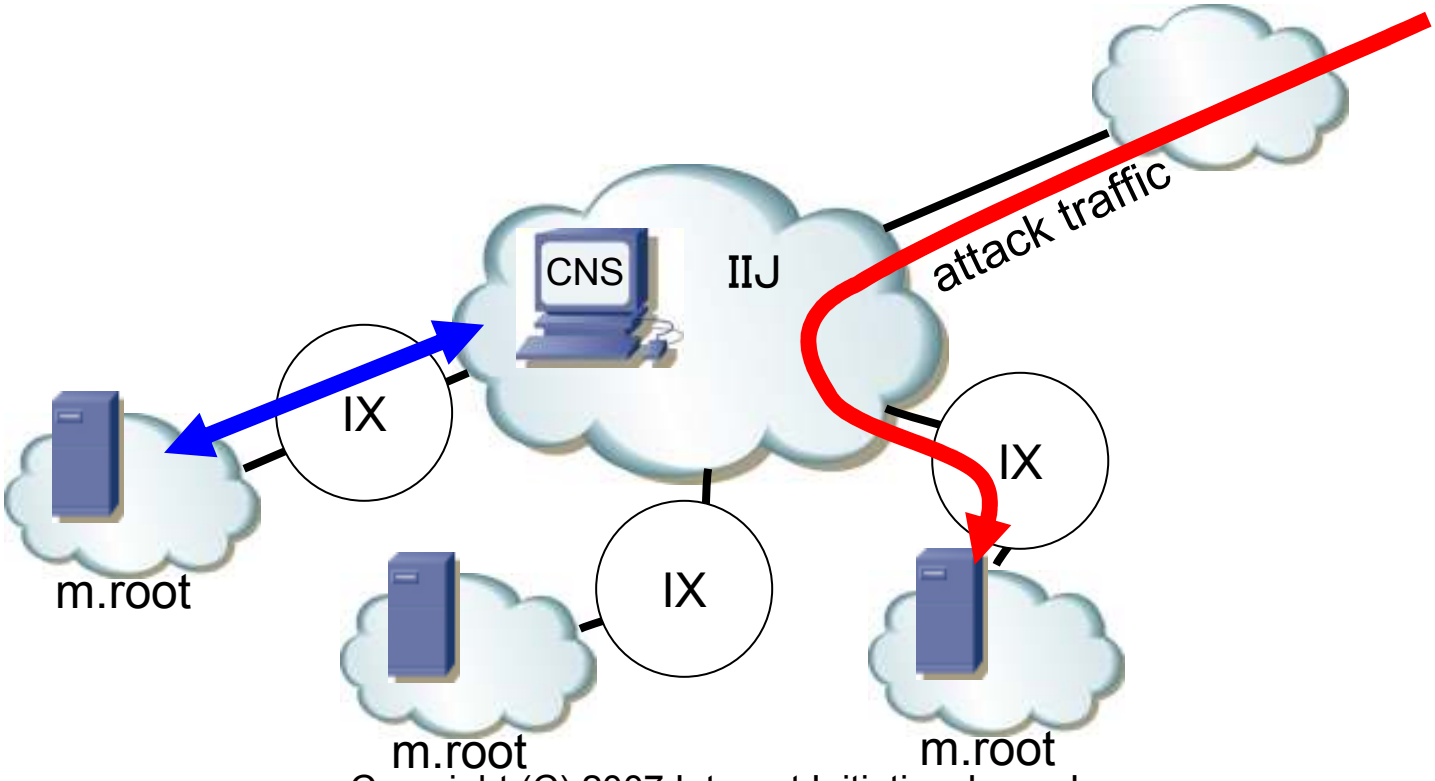

response delay of f.root

hostname.bind. - "kix1b.f.root-servers.org"

during attack

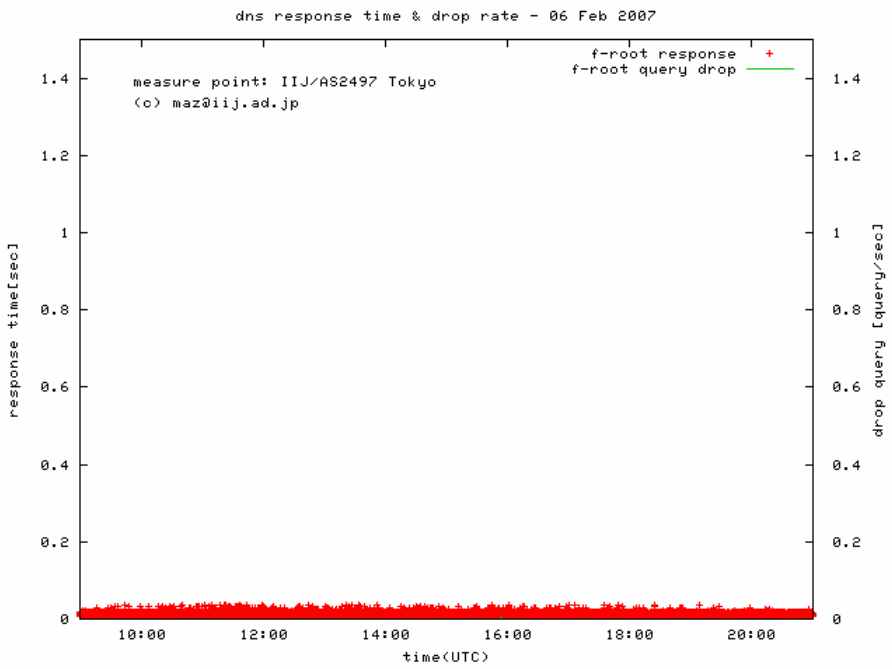

1 week later

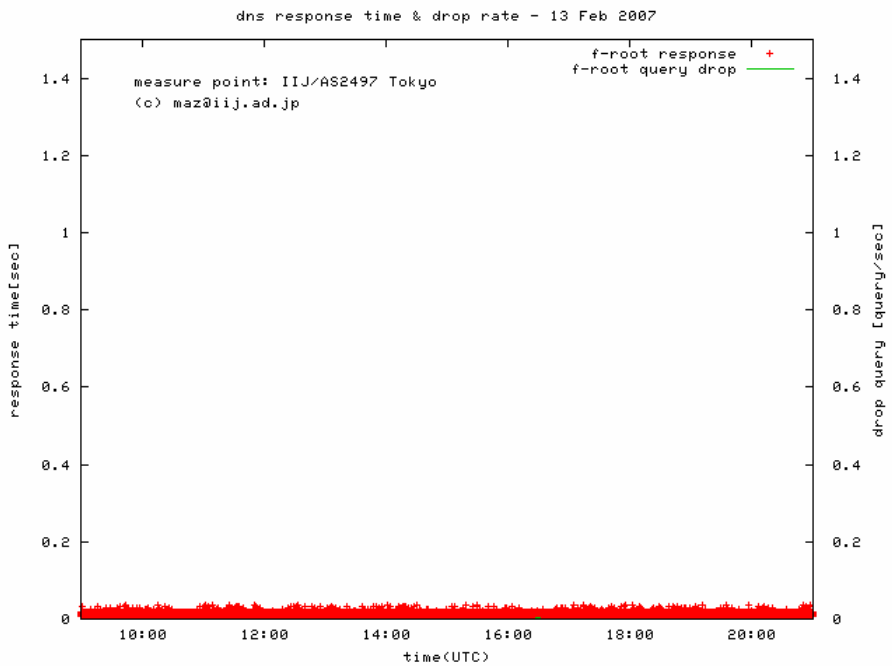

response delay of i.root

hostname.bind. - "s1.tok"

during attack

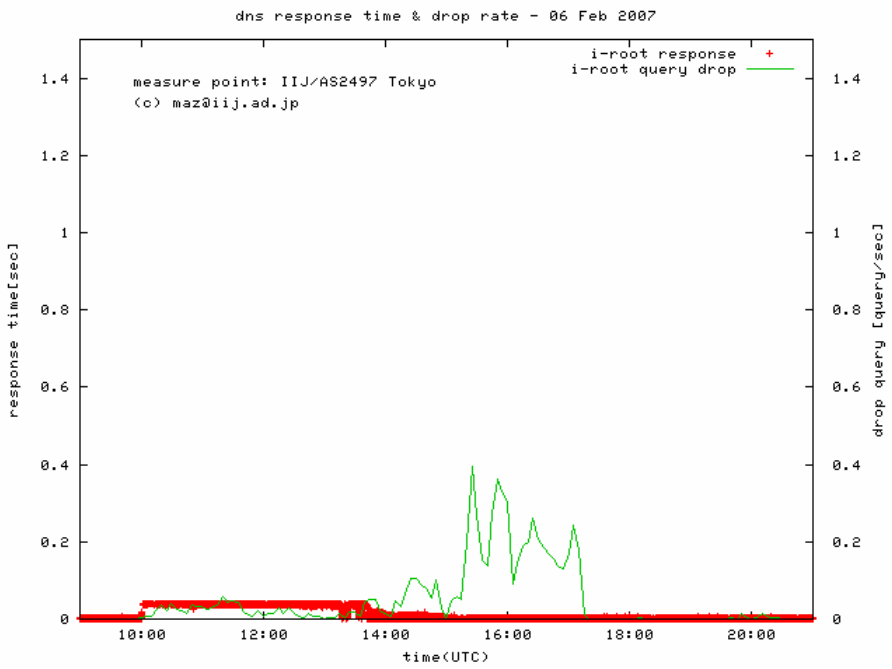

1 week later

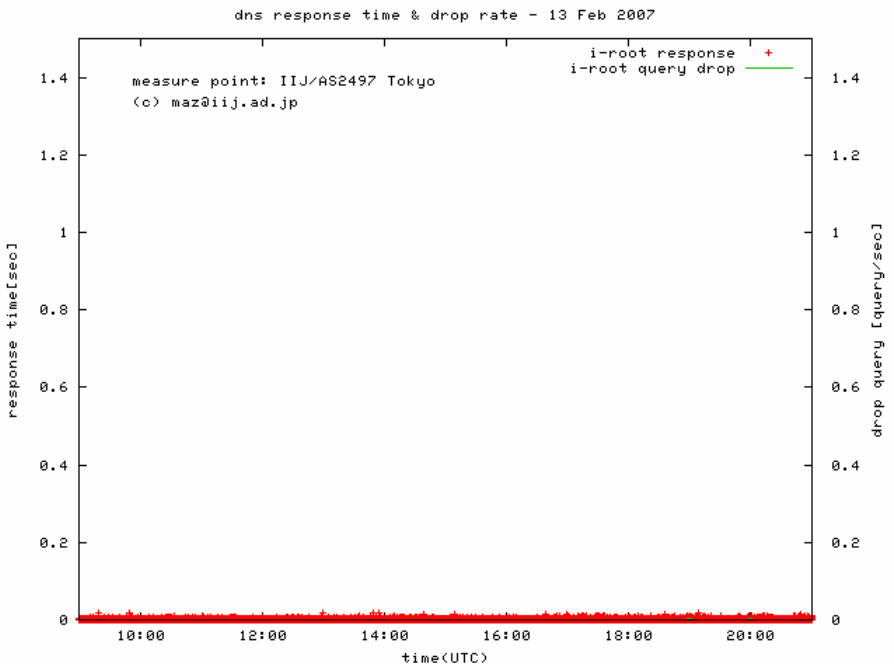

server selection during attack

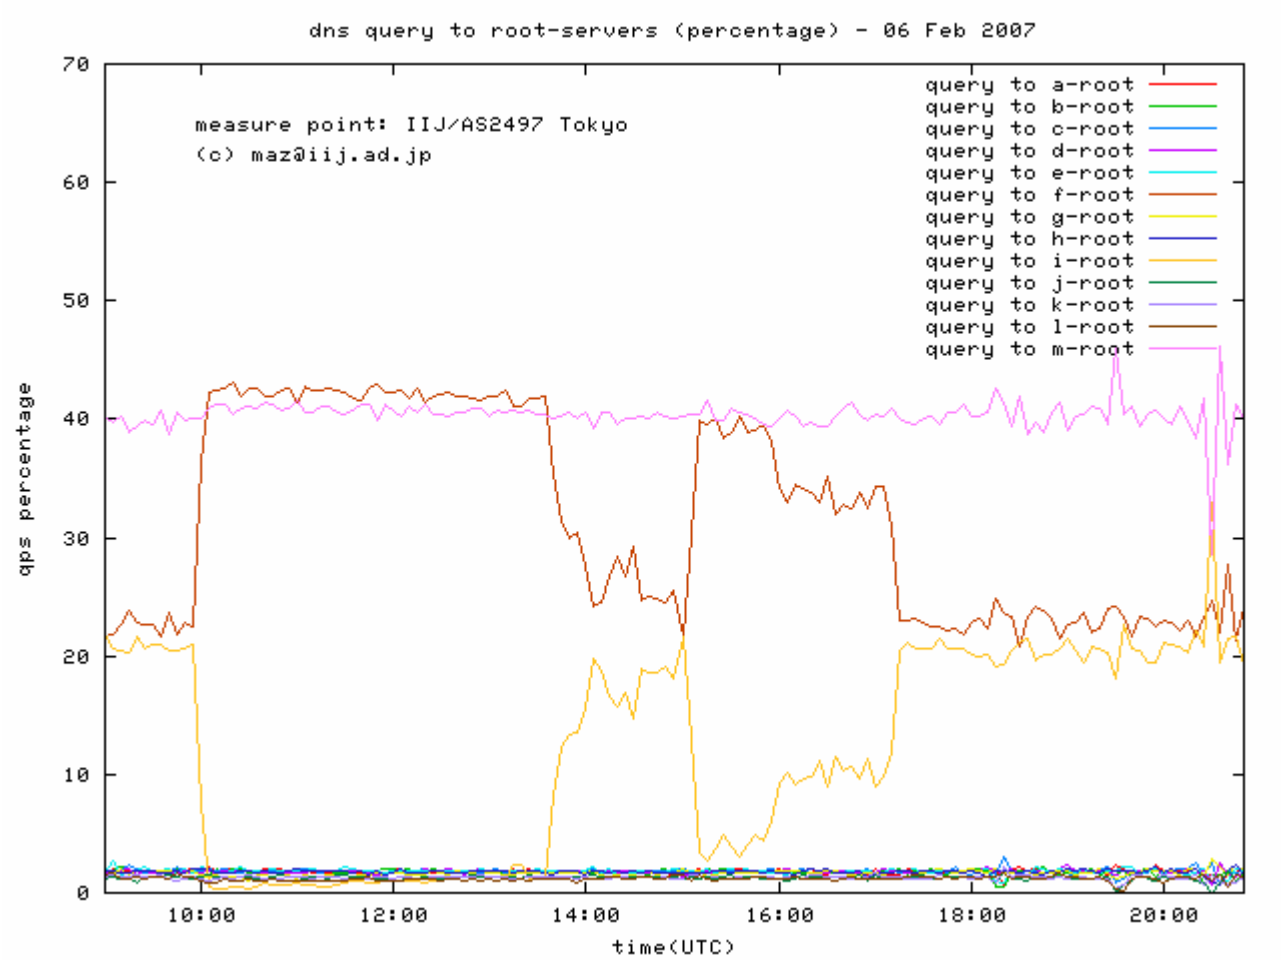

application layer restoration

- DNS Cache servers selects stable authoritative servers automatically.
- Of cause, this feature depends on its implementation.

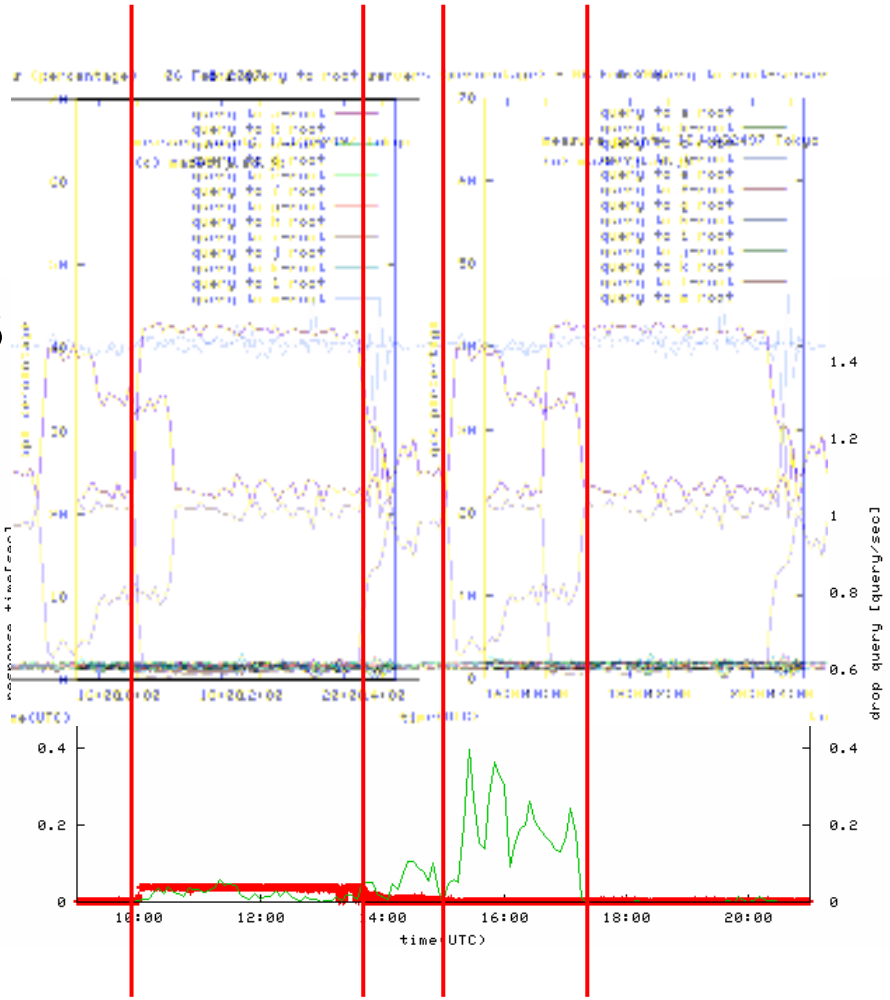

response delay of b.root

hostname.bind. - "b2"

during attack

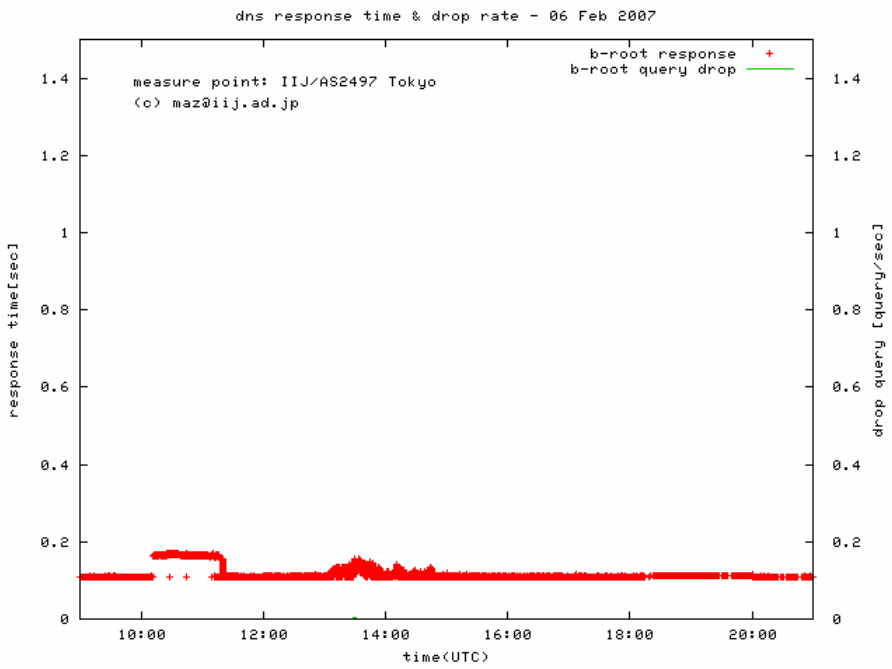

1 week later

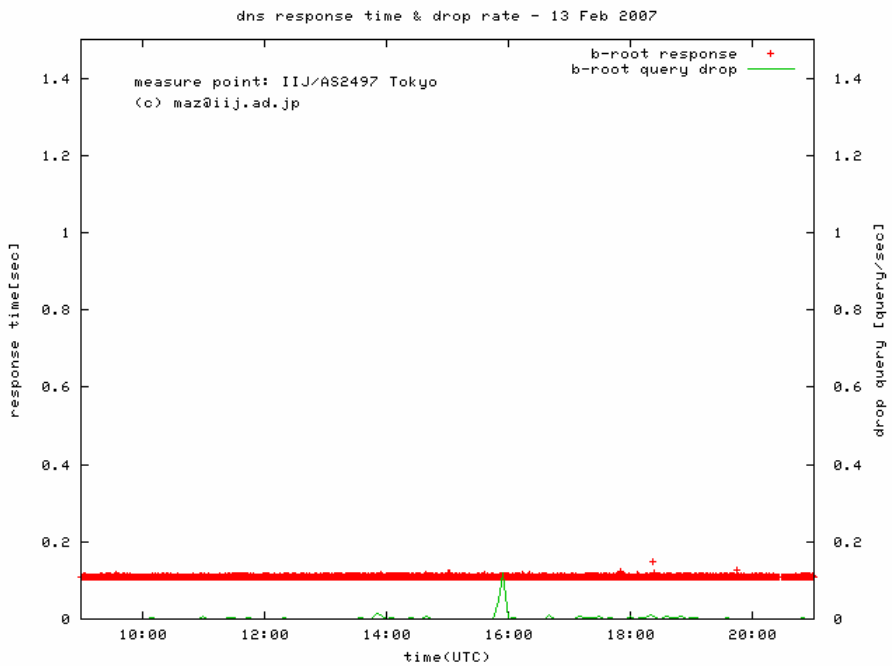

response delay of g.root

hostname.bind. - "g.root-servers2.net"

during attack

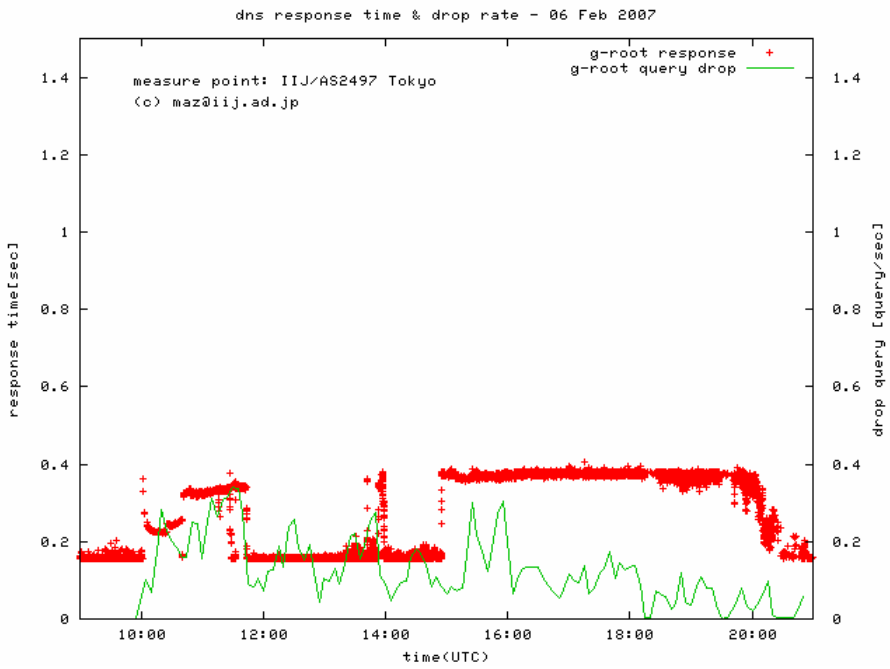

1 week later

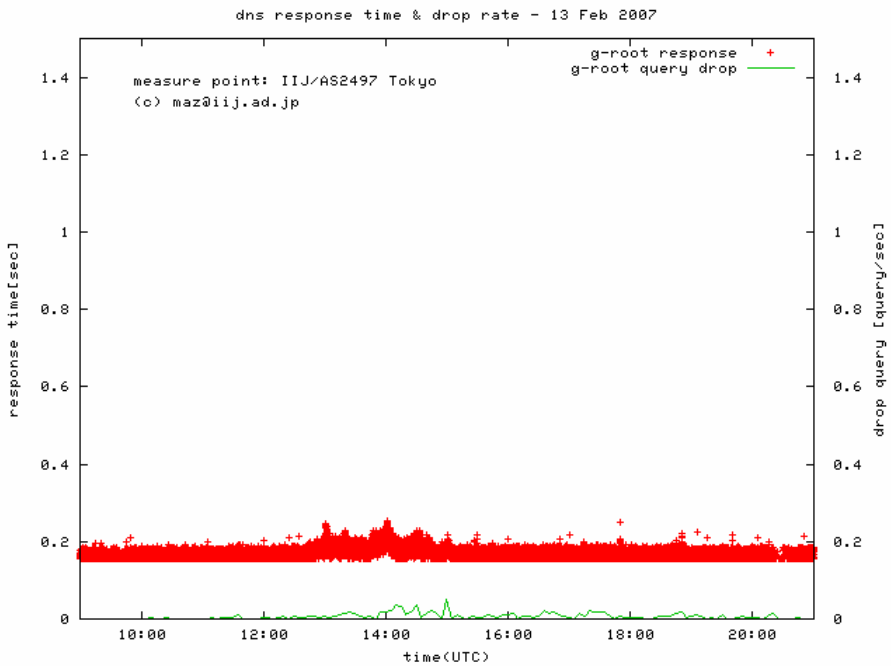

response delay of k.root

hostname.bind. - "k1.linx"

during attack

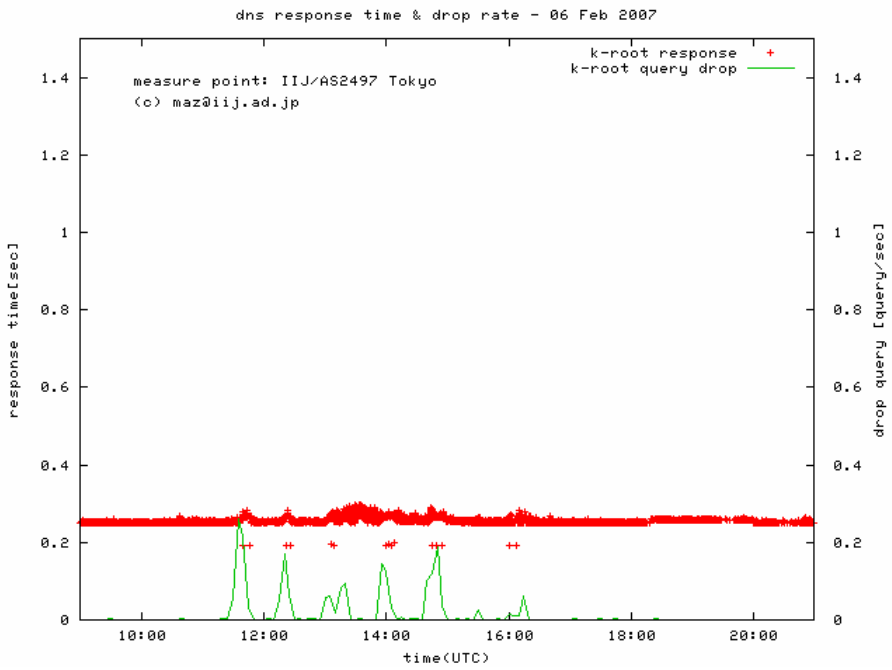

1 week later

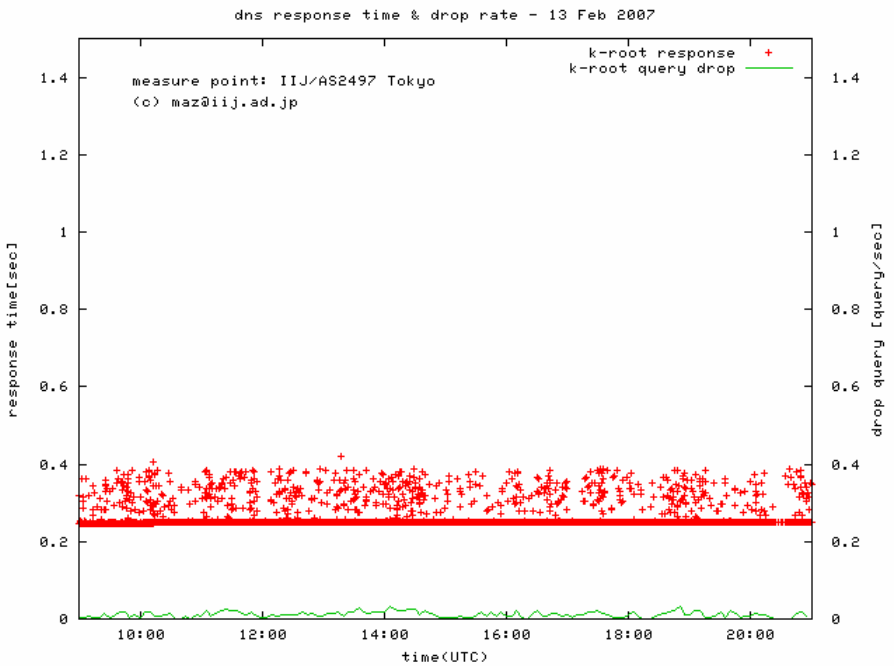

response delay of l.root...?

hostname.bind. - "lax-25"

during attack

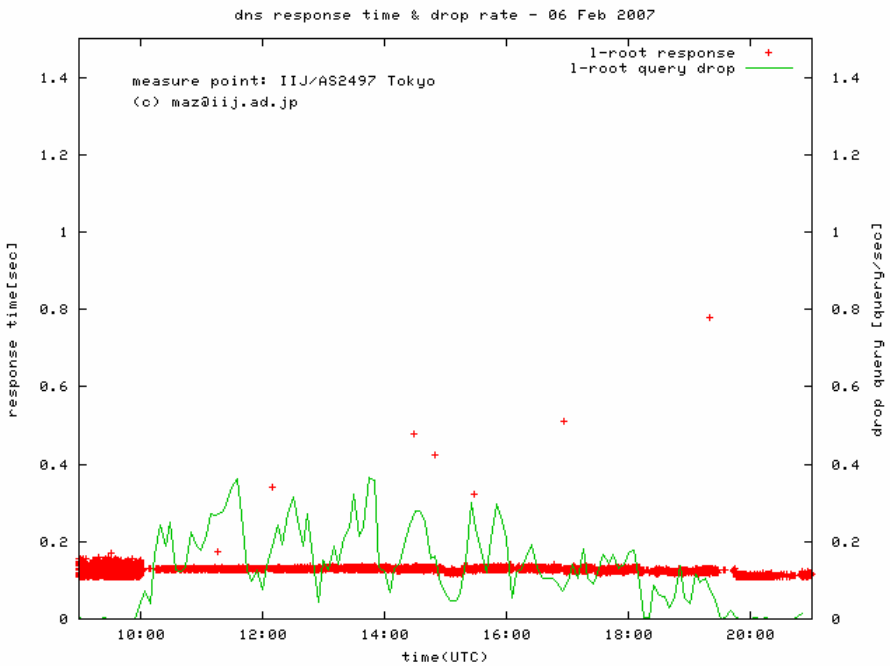

1 week later

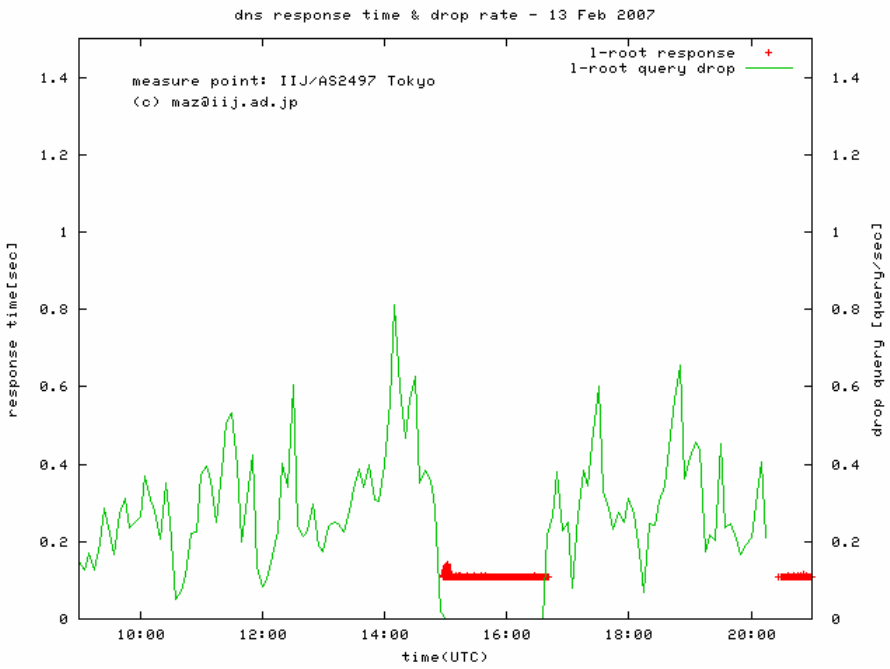

of queries to root-servers

- 1229097 total queries
  - 1223957 invalid_TLD (99.5%)
    - 1110543 AforA (90.3%)
    - 113414 other invalid_TLD (9.2%)
  - 5140 valid_TLD(0.4%)
    - 4787 .arpa (0.3%)
    - 353 other valid_TLD(0.02%)

duration 08 Feb 2007 09:00UTC-21:00UTC

conclusion

- There was a attack, but we can say the effect to end-user is minimal or ignorable.
  - anycast works fine. ☺
  - application layer restoration works fine.☺
  - thanks for the long TTL, cache servers need to send a query to root-servers sparsely.
- But we found delays on .org response, we need further researches about this.