TTL Considerations for Infrastructure Resource Records

Vasilis Pappas, Dan Massey, Lixia Zhang

DNSOP WG 3/19/07

"Development of the Domain Name System"

- "a low TTL is desirable in that it minimizes periods of transient inconsistency, *while a high TTL minimizes traffic and allows caching to mask periods of server unavailability due to network or host problems*."

—Mockapetris, SIGCOMM 1988

Paul made a general statement

our interest is to show the benefit of setting proper TTL value for *infrastructure RRs* (= NS & corresponding A RRs)

Evaluating TTL impact: Methodology

- Experiments:
  - 5 DNS traces from 4 univ. CS depts and 1 national ISP, 7-day long
  - Simulate DNS cache
  - On the last day: assuming *All TLD servers down* for 3, 6, 12, or 24 hours
- Question: How many queries fail?

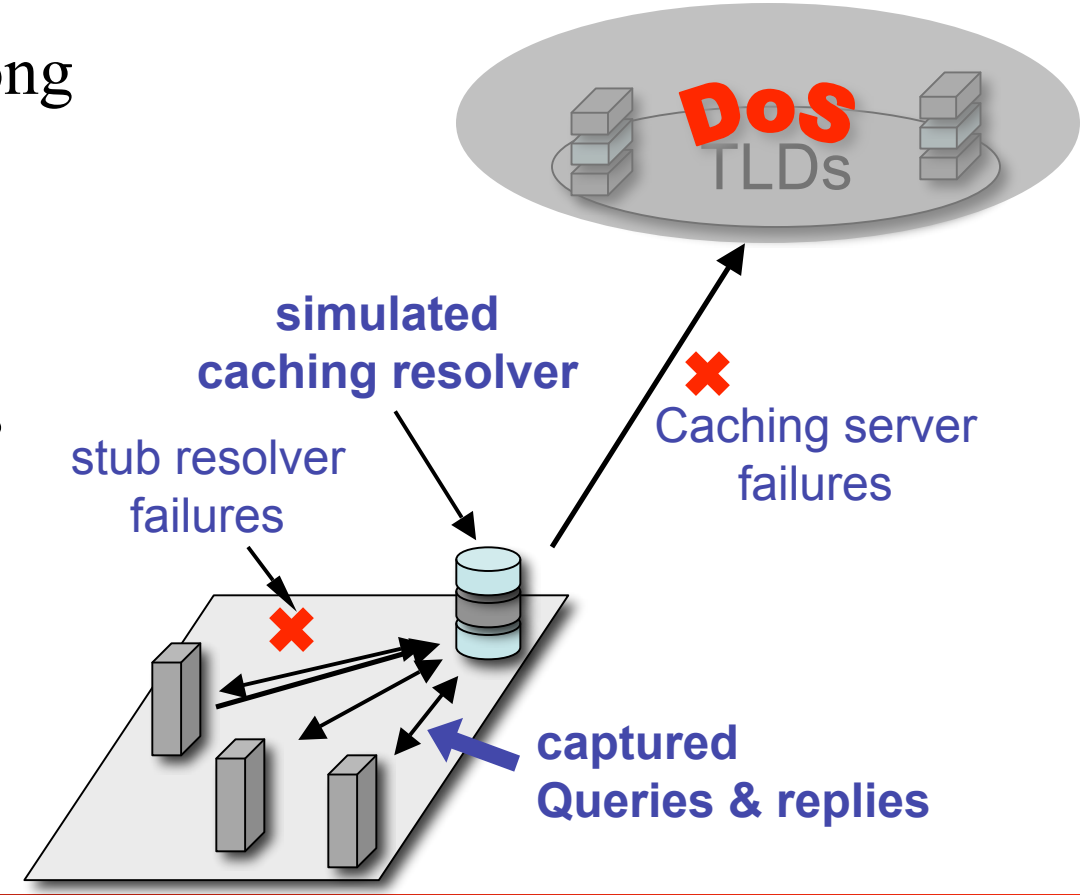

Query failure rate with existing TTL setting

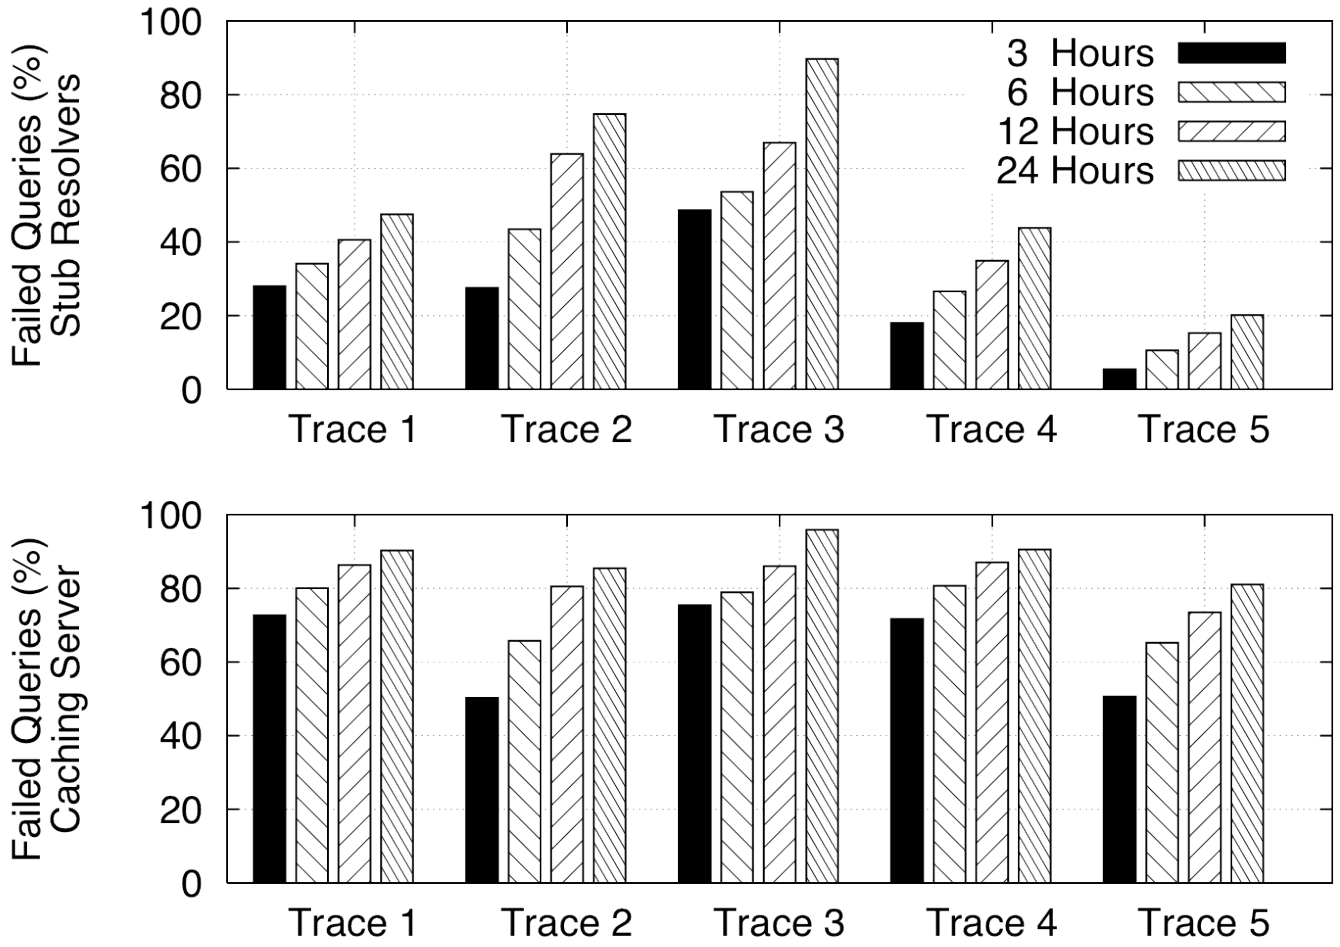

Query failure rate with longer TTL setting

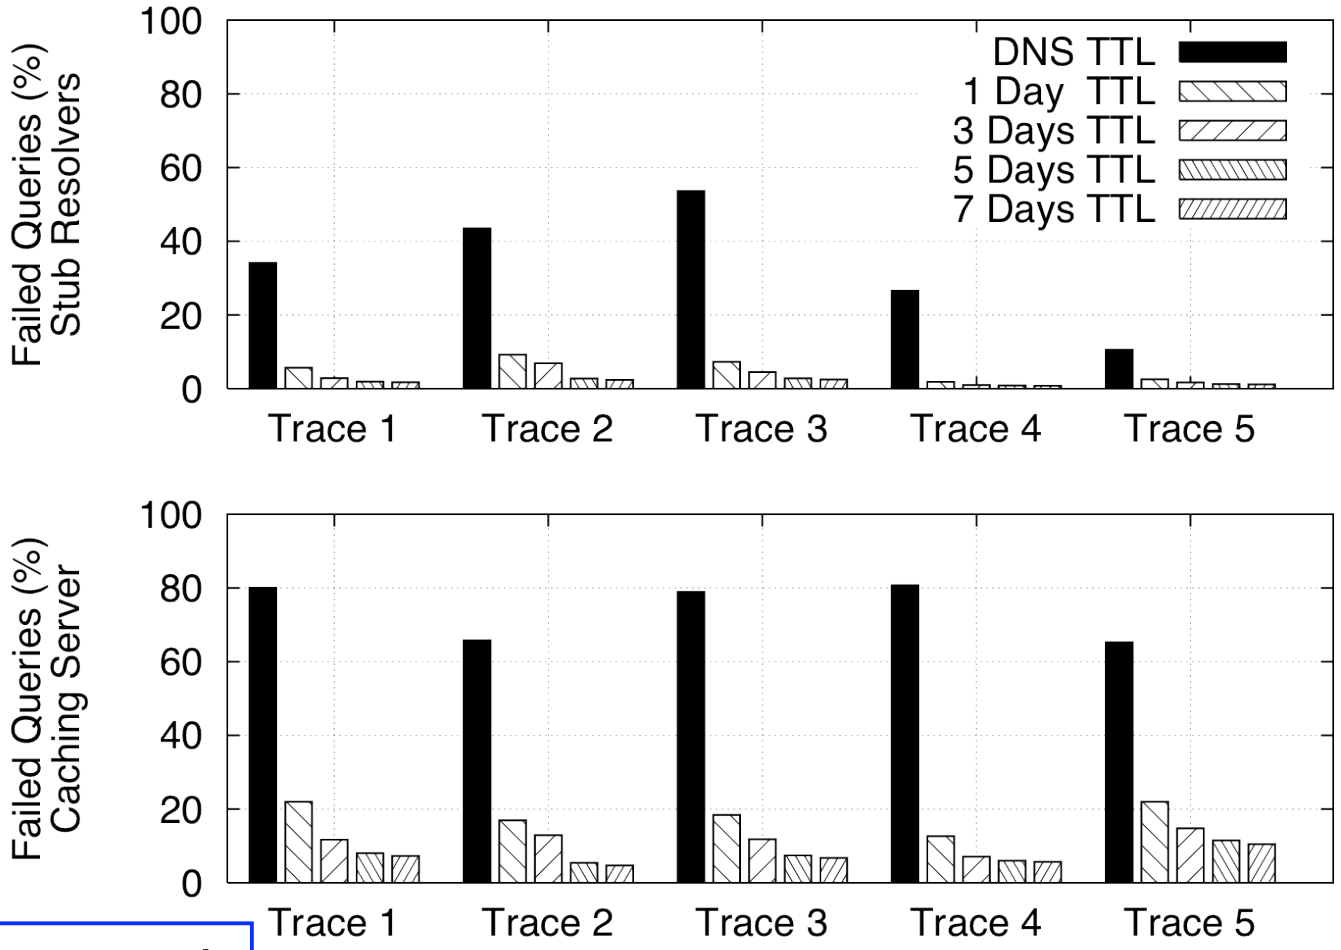

Overhead
(-6)~(-14)%

6-hour attack

The current TTL setting for infrastructure RRs

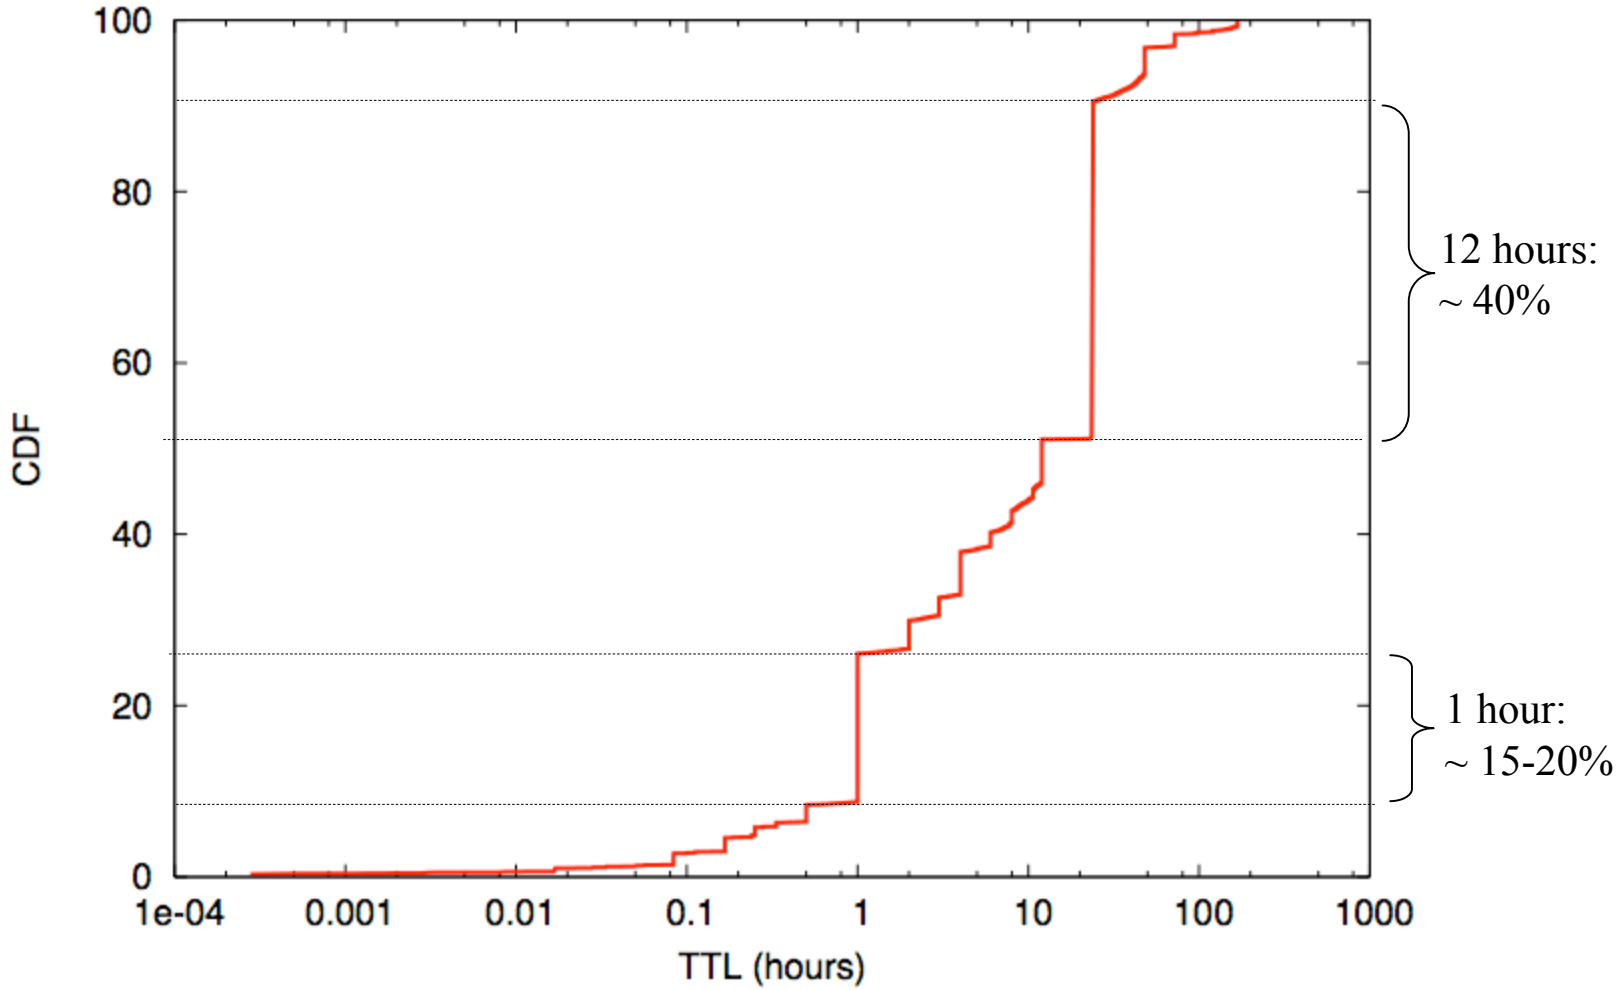

Collected from our traces, note almost 10% RRs have TTL ≤ 1 hour

How frequently do infrastructure RRs change in real life?

- Randomly selected 100,000 zones out of 15 million
  - queried each of the 100,000 zones twice a day to obtain its infrastructure RRsets over a 4-month period
- **75%** of the zones made no change to their NS or corresponding A RRSets during the entire 4 months
- **11%** of the zones showed changes to their NS Rrset
  - **5%** of the changes occurred in less than 2 months
- 22% of the zones changed their servers' A records during the 4 months
  - **10%** of the changes occurred in less than 2 months

What if infrastructure RRs change before their TTL expires

- Basically delayed response, *not* service failure
  - If at least one authoritative server of the zone did not change: query will eventually succeed
    - And bring to the caching resolver the updated NS & A RRs
  - If all authoritative servers of a zone *Z* changed before their TTL expires: the query eventually clamb up the tree and walk down
    - hopefully the changes have made to the parent of Z
- Only one known problematic case so far: when one changes DNS service provider from A to B, the former provider A can cause damage
  - Hopefully this case is unlikely to happen to higher level zones

Benefits from setting longer TTL for infrastructure RRs

- Effective solution for improving DNS availability when higher level servers are unavailable
  - Even a moderate value, e.g. TTL=24hours, can bring significant benefit to availability
- Simple zone config. change; no protocol change
- Any zone can choose to set a longer TTL for their NS+A RRset to improve *their own availability*
  - In contrast to anycast solution which may not be generally affordable

Seeking input from DNS operators

- We plan to write a "TTL considerations for infrastrucdture RRs" draft
- To see the coin from all sides, we need a better understanding on:
  - What are the cases that may need short TTL for infrastructure RRs?
  - How short do the TTLs need to be?