DNS PREFETCHING: WHEN GOOD THINGS GO BAD

Srinivas Krishnan and Fabian Monrose

THE UNIVERSITY
of NORTH CAROLINA
at CHAPEL HILL

Information quest

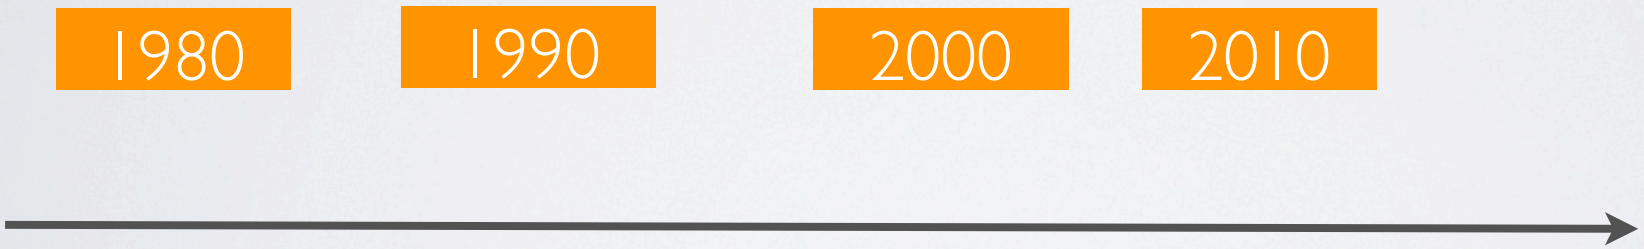

Information quest

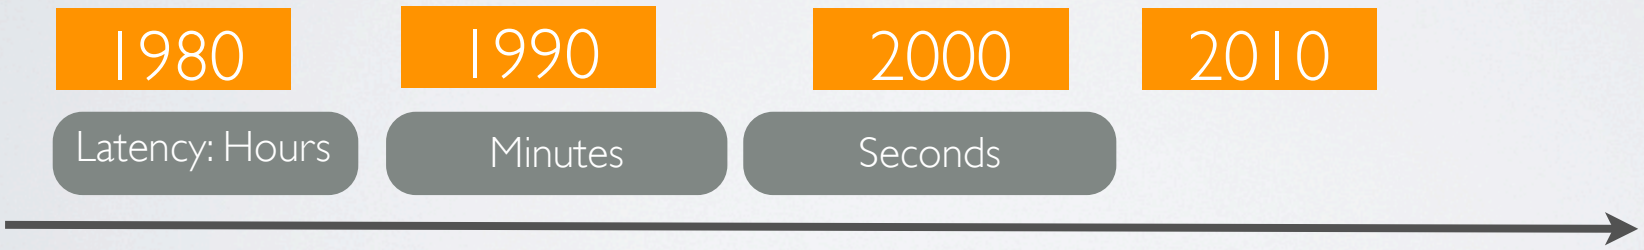

Information quest

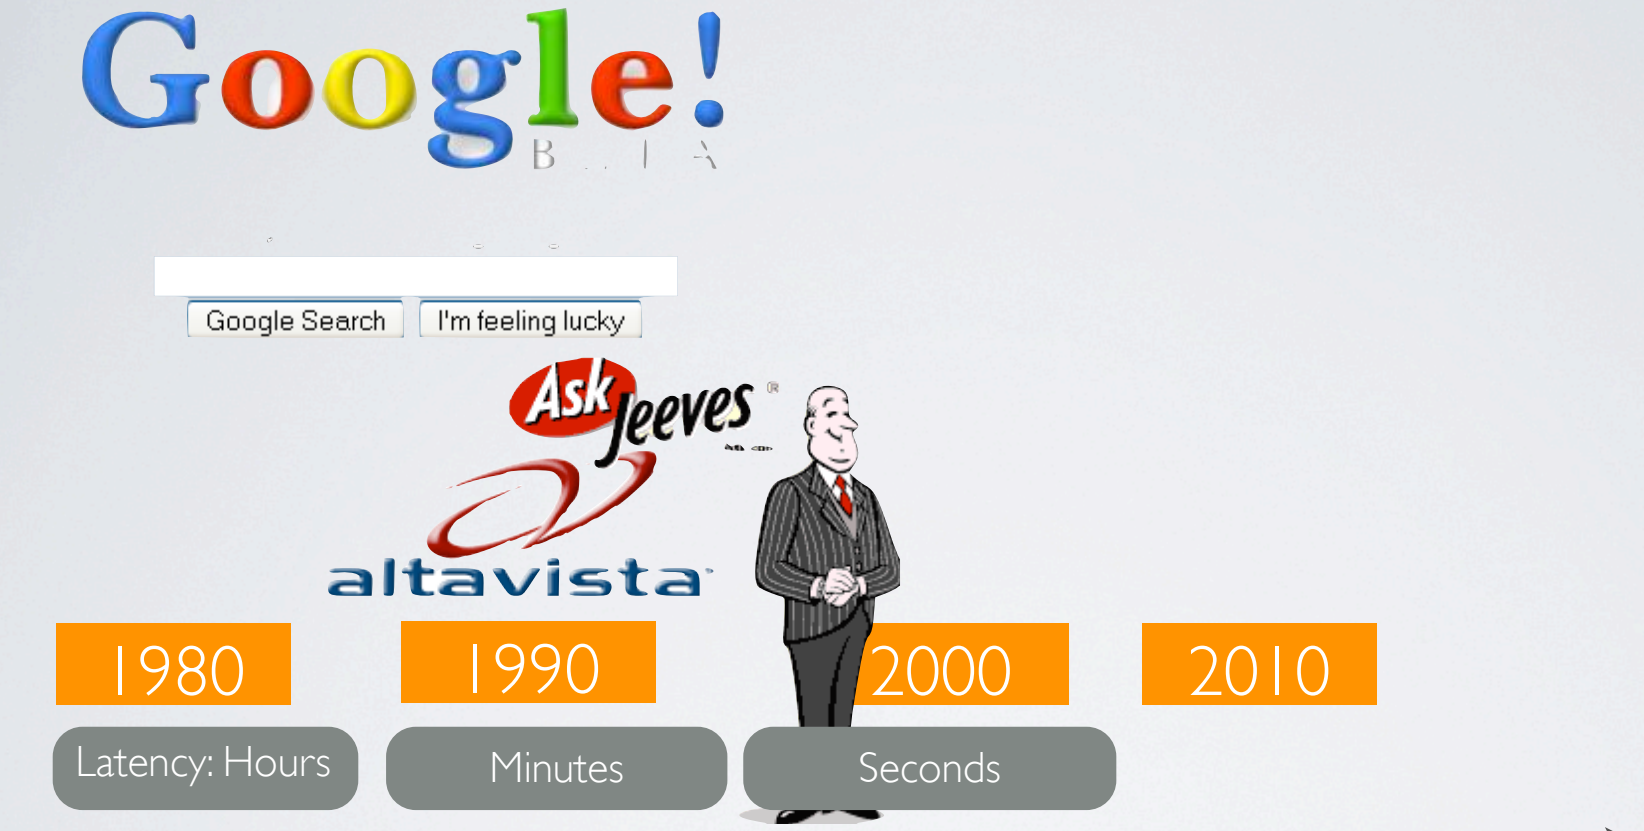

Timeline

Information quest

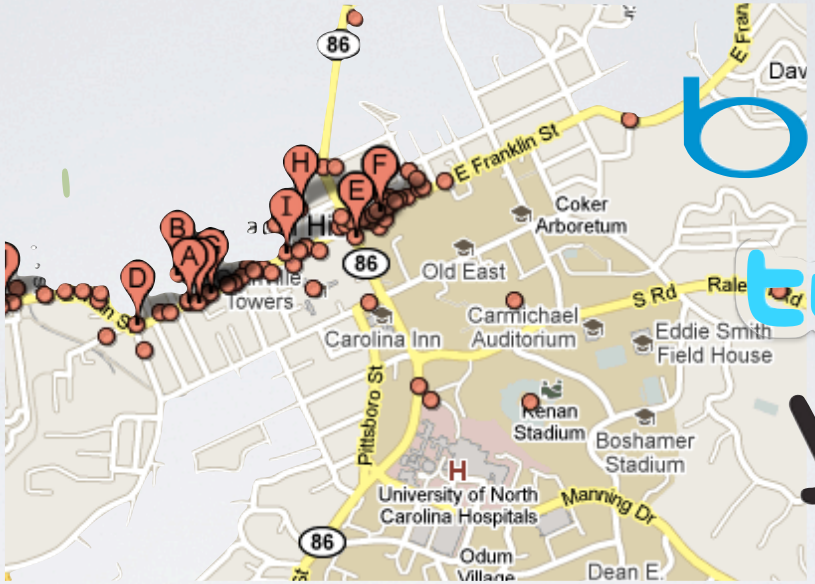

1980 | 1990 | 2000 | 2010

Latency: Hours | Minutes | Seconds | Milliseconds

Timeline

2

Browser Wars

Scripting

Render

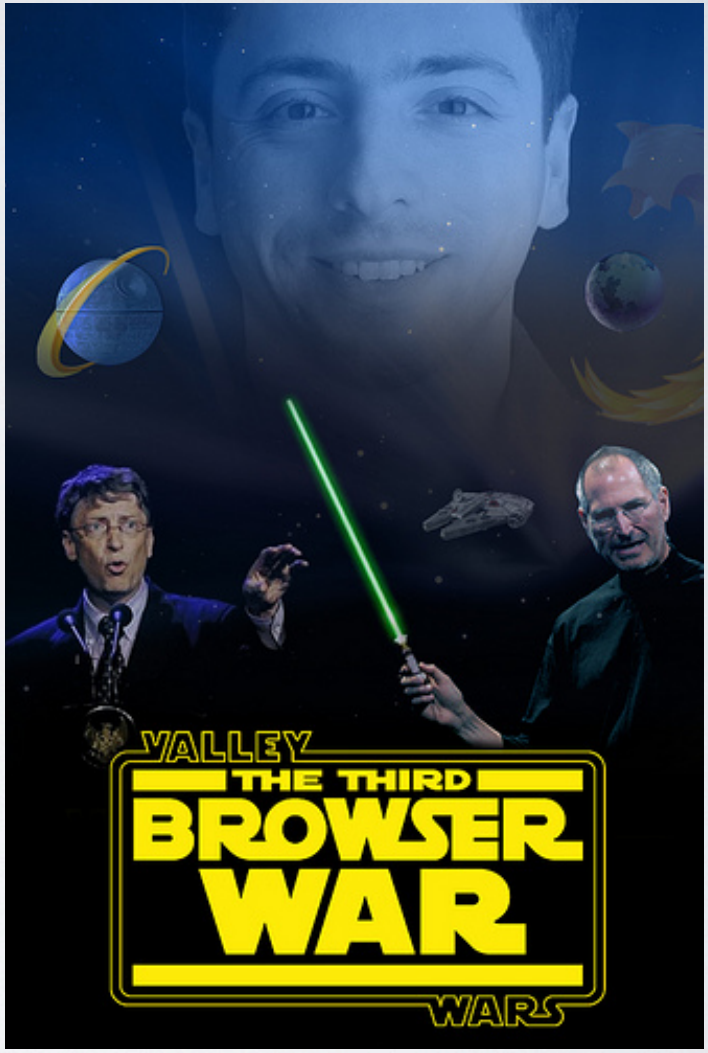

Browsing and DNS

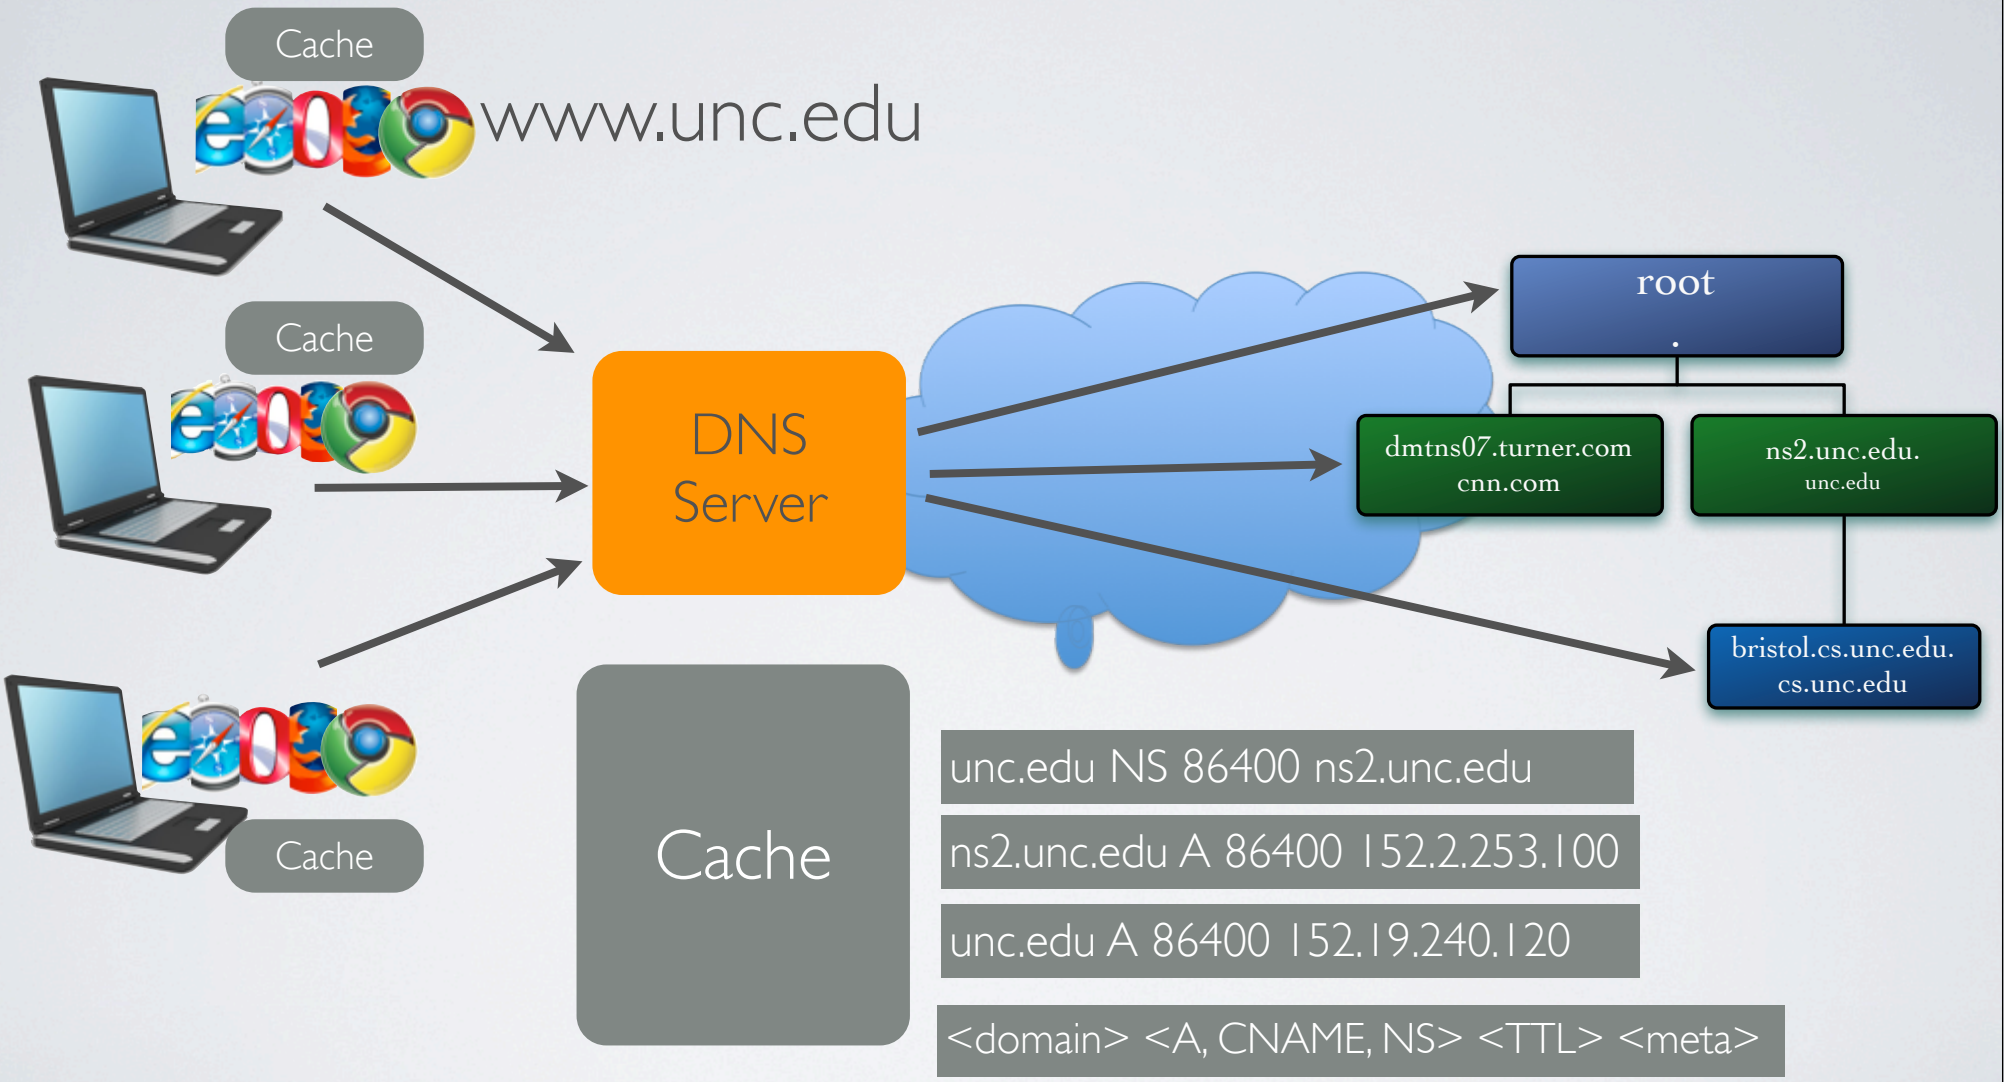

DNS Optimization

- Proactive DNS pre-resolutions
- Two basic approaches:
  - Guess as the user types
  - Fetch <href> links from a rendered page
- Focus on reducing user perceived latency

5

DNS PRE-RESOLUTION

Gambling Addiction

DNS Server

Cache

www.google.com CNAME 586186 www.l.google.com

www.l.google.com A 60 www.l.google.com

6

DNS PRE-RESOLUTION

Gambling Addiction

Google Search | I'm Feeling Lucky

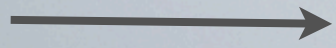

DNS Server

www.google.com CNAME 586186 www.l.google.com

www.l.google.com A 60 www.l.google.com

sac.edu

Cache

6

DNS PRE-RESOLUTION

Gambling Addiction

Google Search | I'm Feeling Lucky

sac.edu

DNS Server

Cache

www.google.com CNAME 586186 www.l.google.com

www.l.google.com A 60 www.l.google.com

sac.edu A 73136

6

DNS PRE-RESOLUTION

Gambling Addiction

Google Search | I'm Feeling Lucky

sac.edu

DNS Server

Cache

www.google.com CNAME 586186 www.l.google.com

www.l.google.com A 60 www.l.google.com

sac.edu A 73136

gamblersanonymous.org. A 73416
casinogambling.about.com.CNAME 900
treatment-centers.net. CNAME 3600
robertperkinson.com. A 86400
en.wikipedia.org. CNAME 1052
ncpgambling.org. A 73416,
helpguide.org. A 73340
gamblingaddiction.org. A 3600

Prefetching

Privacy Threat

- Reconnaissance of an enterprise
- Ability to track users
- Exploit:
  - Ability to probe a DNS server to infer cache hits.
  - Online probes with target search
  - Offline probe with no prior knowledge

7

Online Probing

Was a target search performed by a client ?

- Build a profile of target search
- Use cache snooping
- Check for presence of profile
- Report

8

Building a Profile

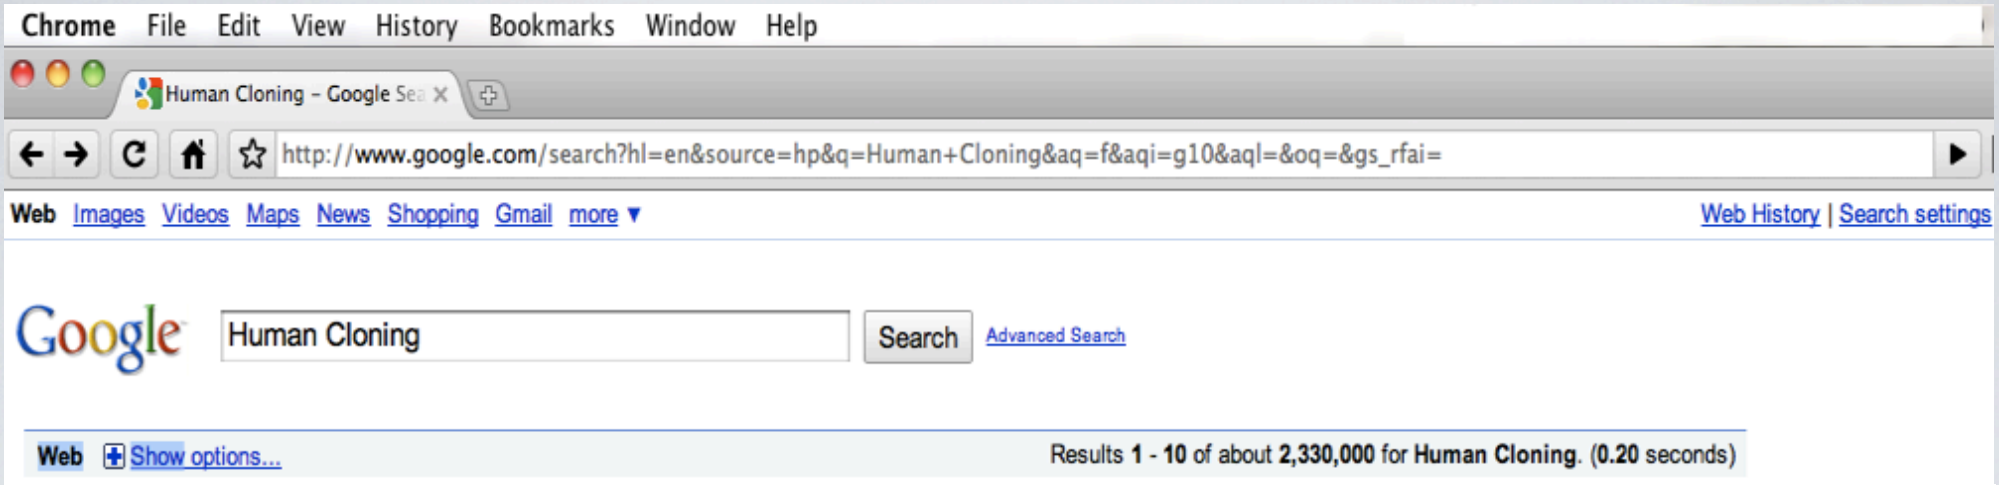

www.howstuffworks.com.

ama-assn.org

learn.genetics.utah.edu.

www.humancloning.org.

www.time.com.

www.ornl.gov.

en.wikipedia.org

www.globalchange.com

www.ncsl.org

Building a Profile

Domains

MinTTL

Decay Curve

howstuffworks.com.
ama-assn.org
genetics.utah.edu.
humancloning.org.
time.com.
ornl.gov.
en.wikipedia.org
globalchange.com
ncsl.org

10

Building a Profile

Domains

MinTTL

Decay Curve

ama-assn.org.
genetics.utah.edu.
humancloning.org.
ornl.gov.
globalchange.com
ncsl.org

10

Building a Profile

Domains

ama-assn.org.
genetics.utah.edu.
humancloning.org.
ornl.gov.
globalchange.com
ncsl.org

MinTTL

ama-assn.org	1800
genetics.utah.edu.	3600
humancloning.org	3600
ornl.gov	86400
globalchange.com	600
ncsl.org	86400

Decay Curve

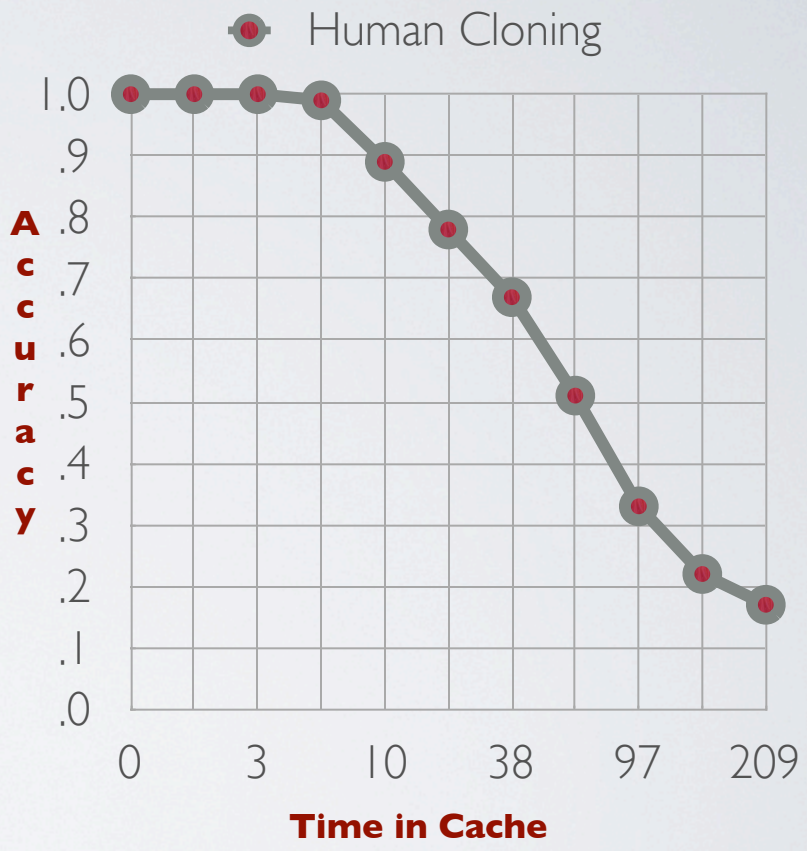

Building a Profile

Decay Curve

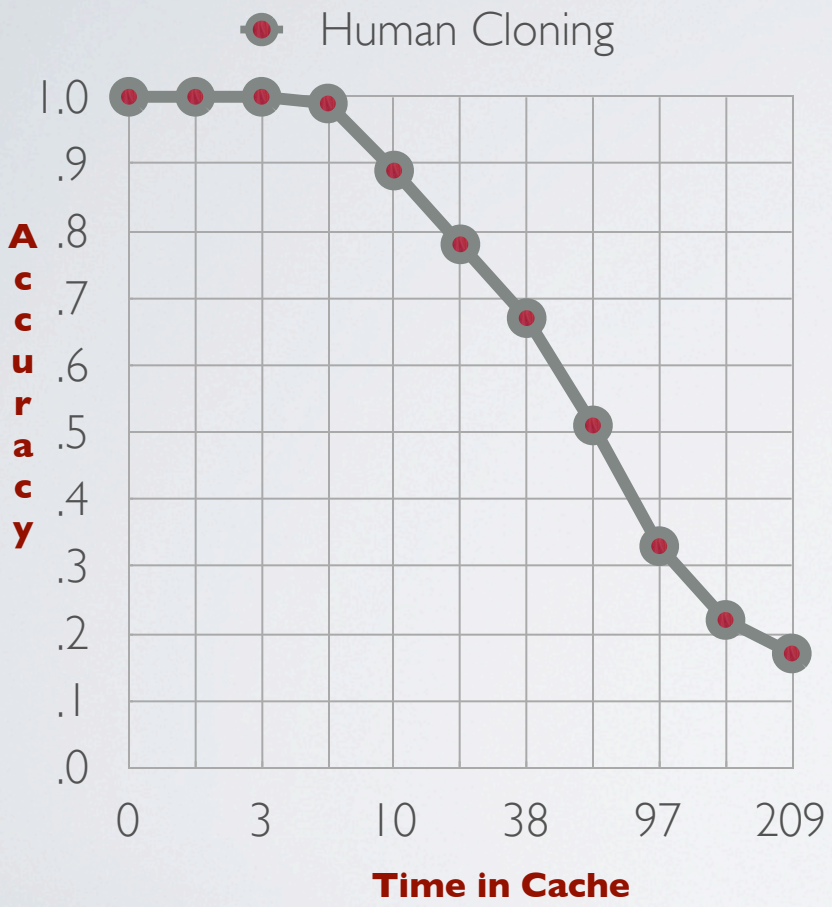

Get Scan Rate

95%	5 Mins
90%	10 Mins
80%	20 Mins
75%	30 Mins
50%	60 Mins

Probe

Attacker

ama-assn.org.
genetics.utah.edu.
humancloning.org.
ornl.gov.
globalchange.com
ncsl.org

genetics.utah.edu ?

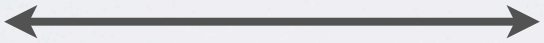

Cache Hit

DNS Server

12

Probe

Attacker

ama-assn.org.
genetics.utah.edu.
humancloning.org.
ornl.gov.
globalchange.com
ncsl.org

ama-assn.org. ?
genetics.utah.edu.?
humancloning.org. ?
ornl.gov.
globalchange.com ?
ncsl.org ?

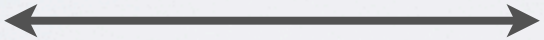

DNS Server

12

Probe

Confidence = % of Elements with same age

Domain	Current TTL	Auth TTL	Age
ama-assn.org	1498	1800	302
genetics.utah.edu.	3298	3600	302
humancloning.org	3301	3600	299
ornl.gov	86099	86400	301
globalchange.com	298	600	302
ncsl.org	86101	86400	299

And if we had access to logs ?

- Can we extract all searches ?

13

DNS Cache: privacy leaks

Goal: Reconstruct Search Term from DNS Cache

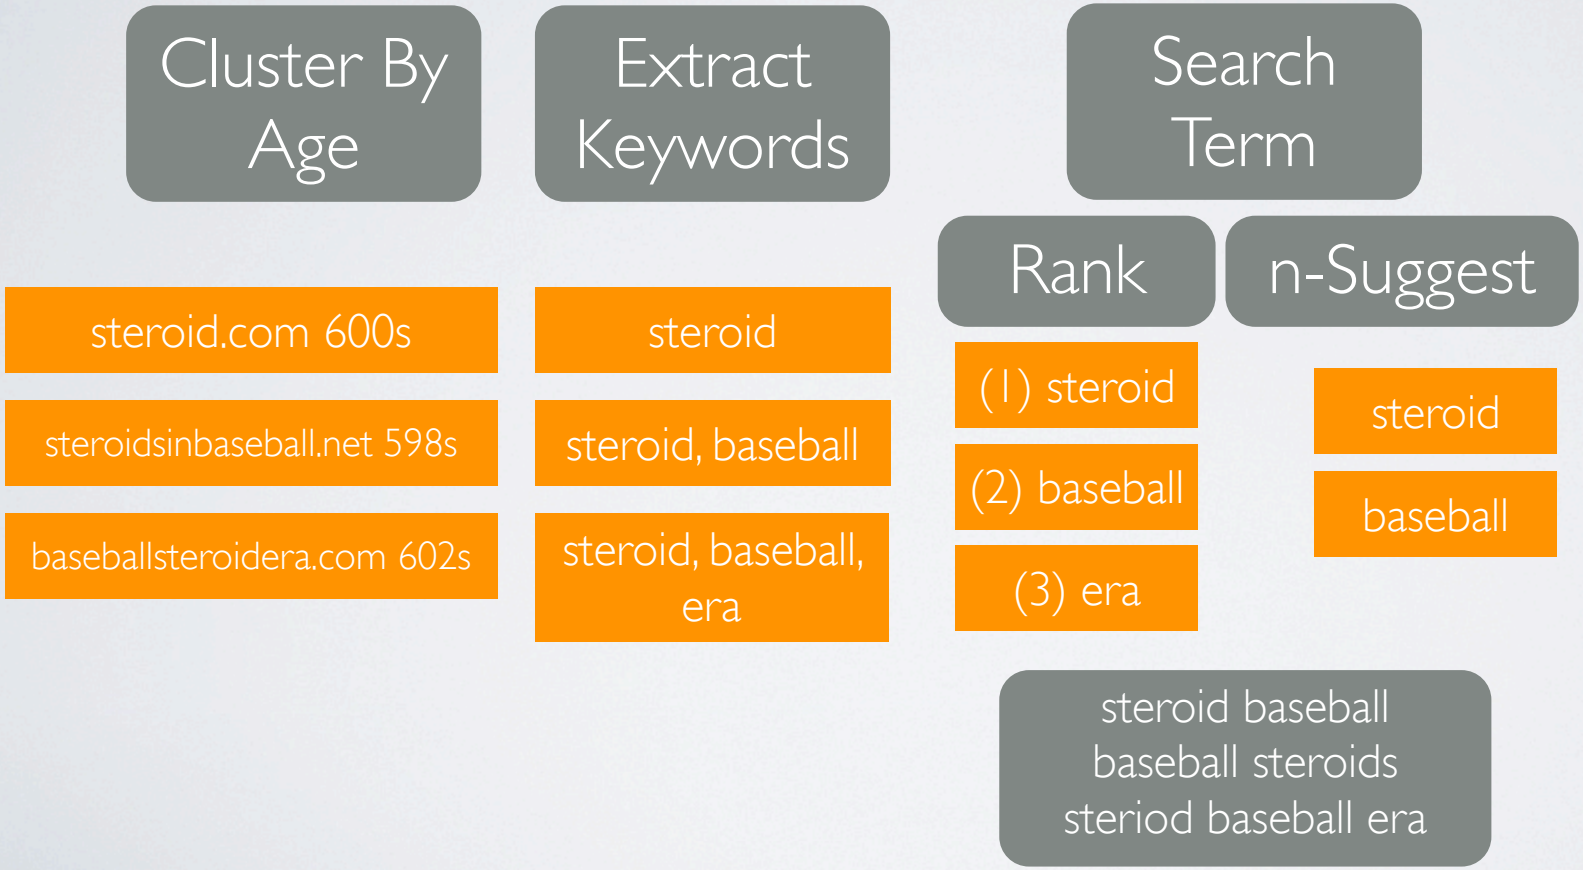

Case I: Preliminary Results

~500 Clients

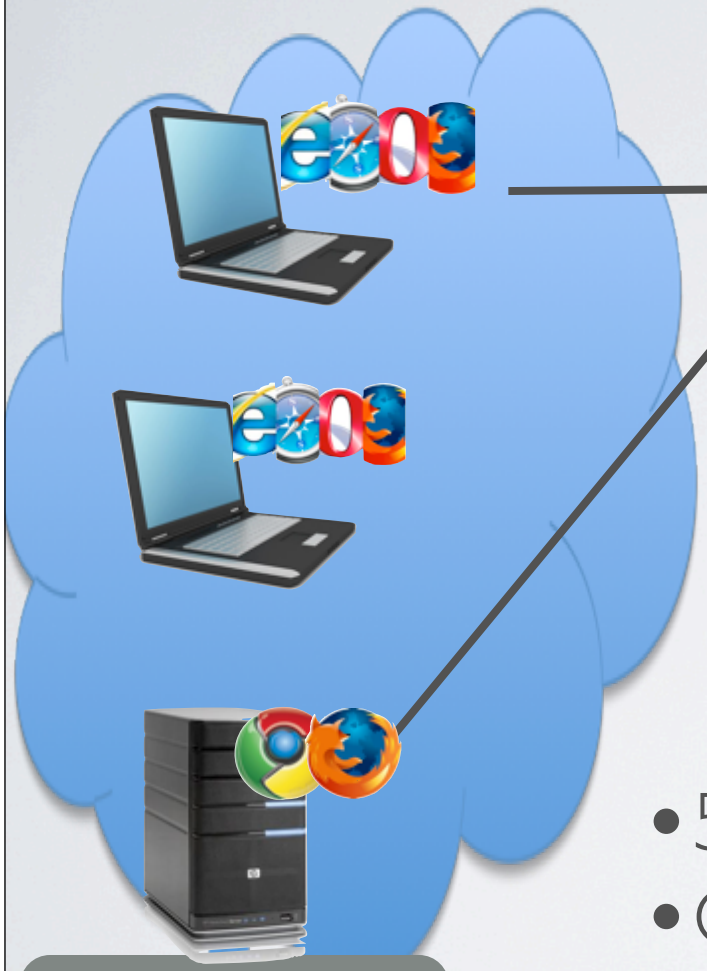

Target DNS Server

Control DNS Server

Inject Queries

- 50 queries
- Over 4 hours
- Variable scan rate

15

Case I: Preliminary Results

~500 Clients

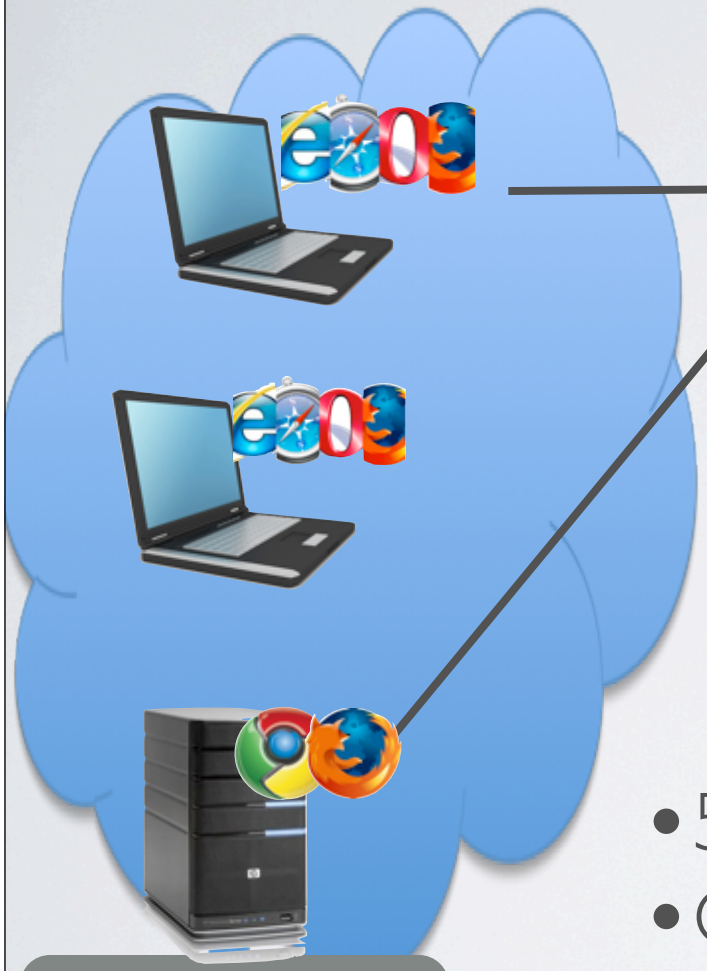

Target DNS Server

Control DNS Server

Build Profile

Inject Queries

- 50 queries
- Over 4 hours
- Variable scan rate

15

Case I: Preliminary Results

~500 Clients

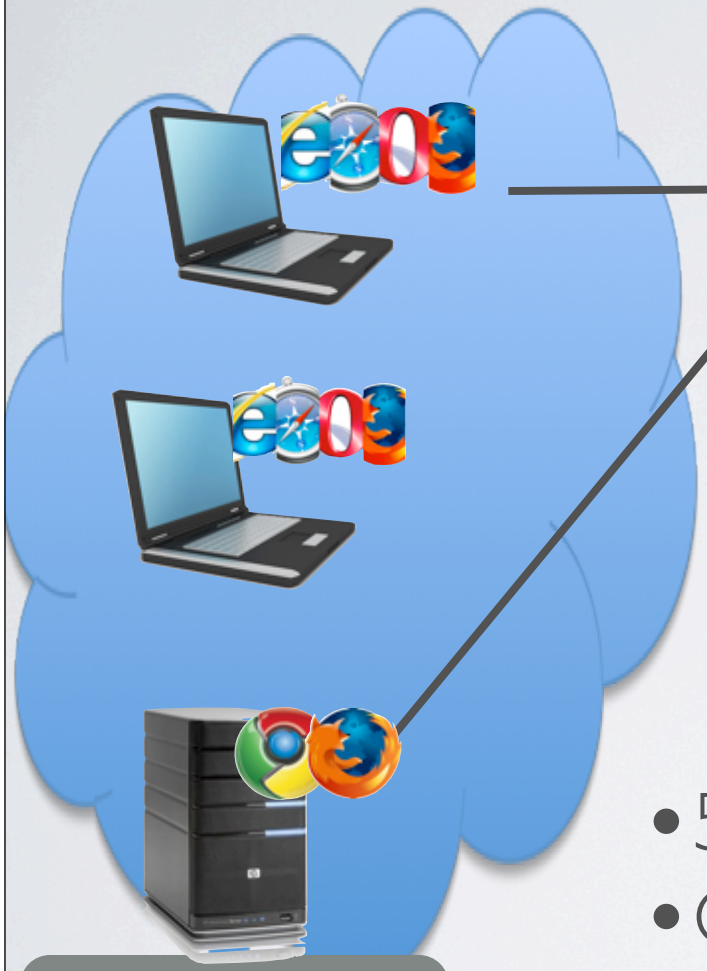

Target DNS Server

Control DNS Server

Probe Server @Scan Rate

Inject Queries

- 50 queries
- Over 4 hours
- Variable scan rate

15

Selected Results

Scan Rate	Average Accuracy
10 Mins	90%
30 Mins	85%
60 Mins	65%

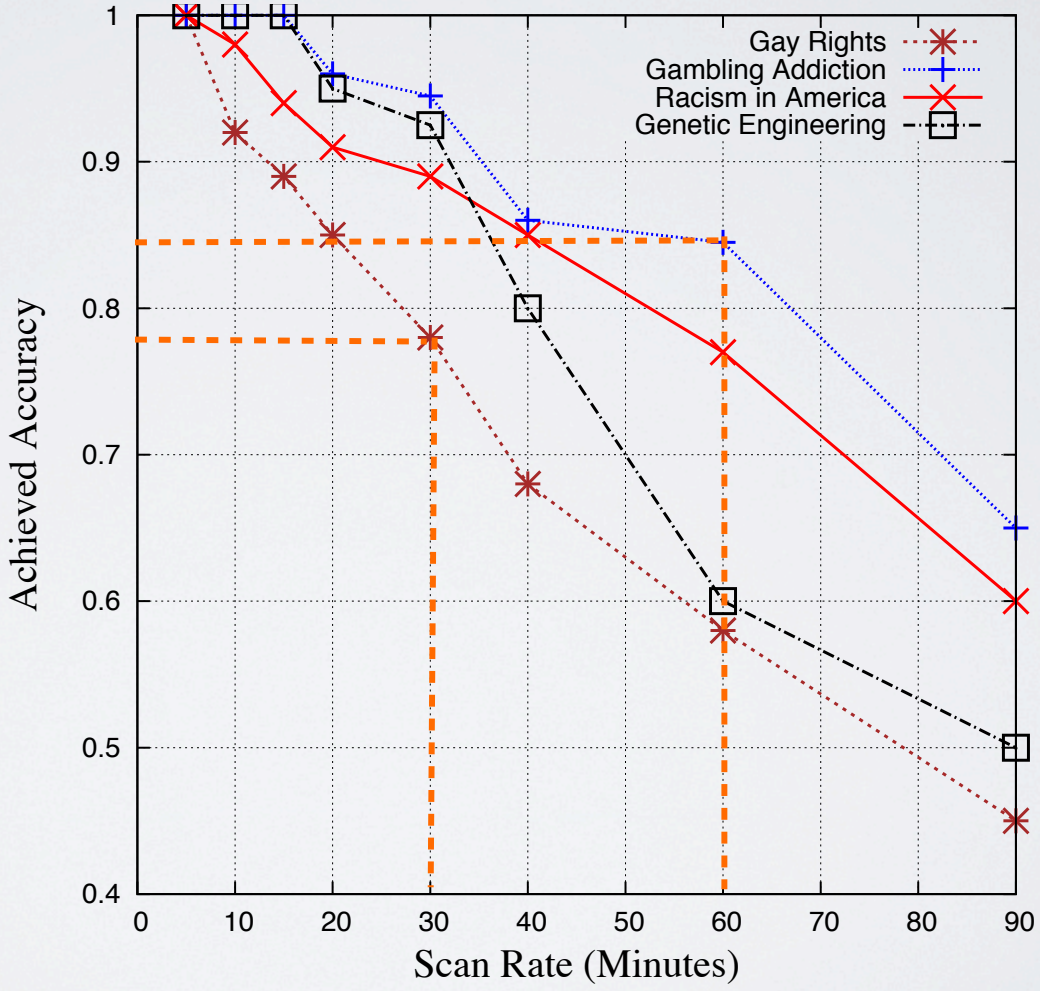

Case II: Preliminary results

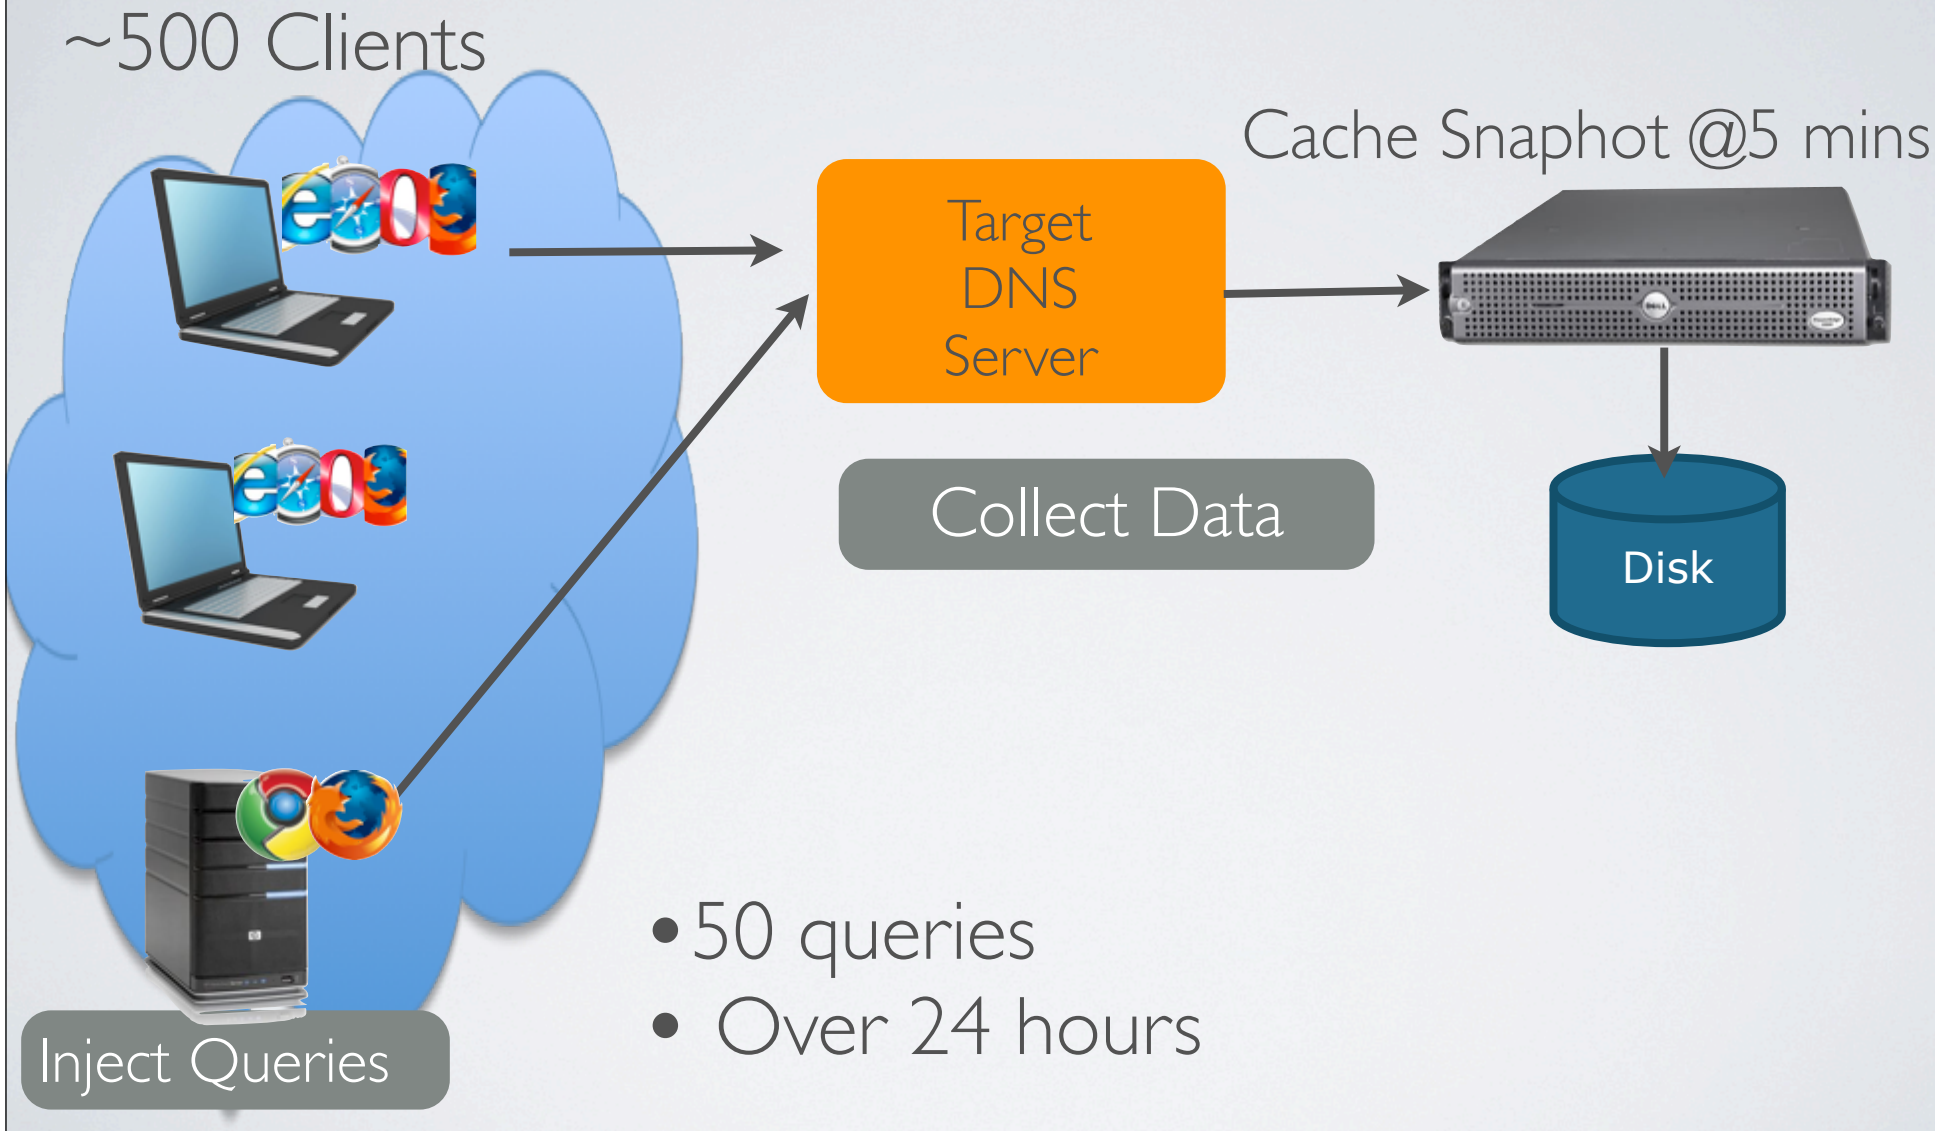

- 50 queries
- Over 24 hours

17

Case II: Preliminary results

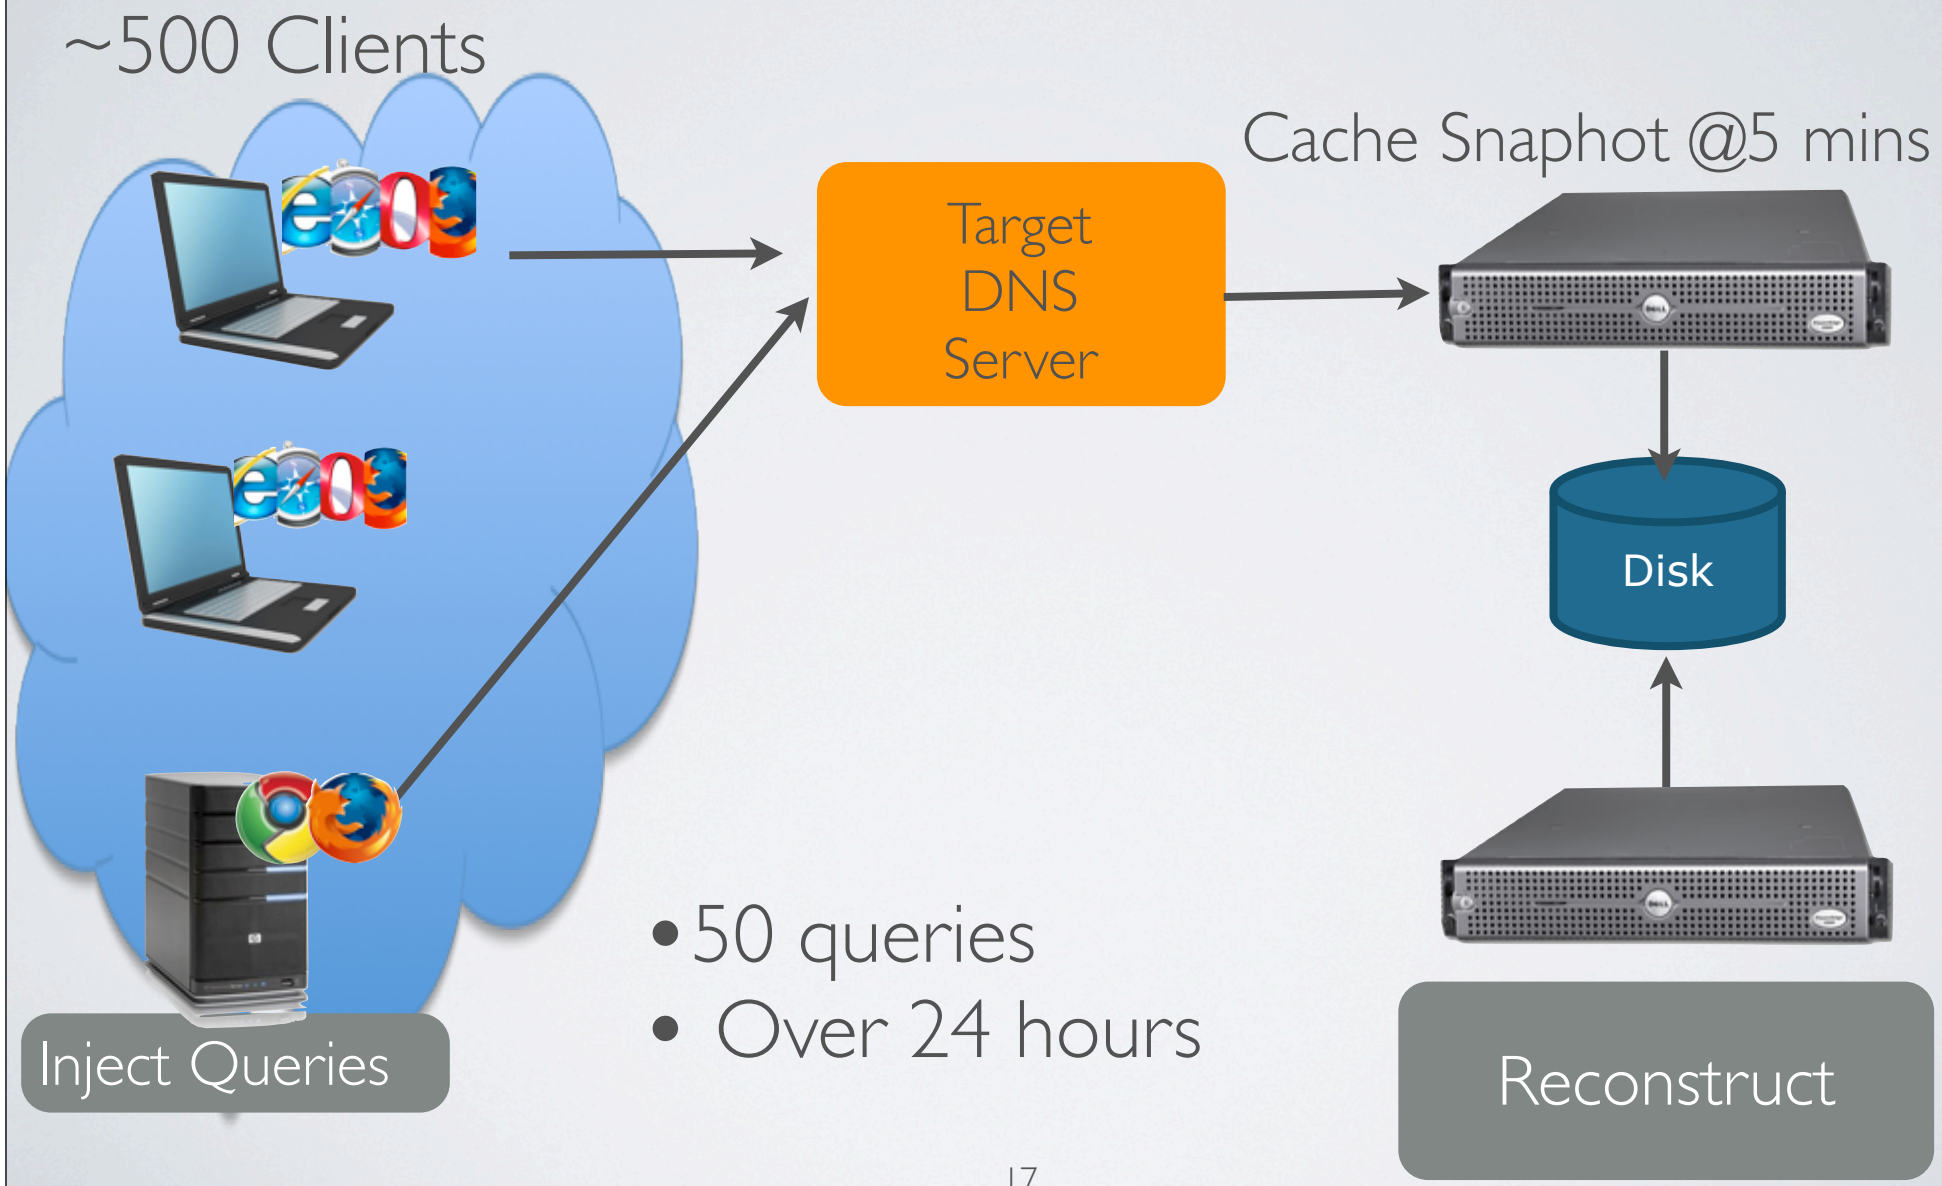

- 50 queries
- Over 24 hours

17

Snapshot of Results

Actual Query	First Guess	Second Guess	Third Guess
Gambling Addiction	gambling addiction	gambling age	addict
Alcohol Withdrawal Syndrome	alcohol withdrawal symptoms	alcoholics anonymous	alcohol poisoning
Gun Control	gunbroker	guns for sale	-
Racism In America	racism america	racism today	racism facts
Biological Weapons	biological warfare	weapons	-

18

Limitations

- Current profiles are non-adaptive, hence searches on "hot topics" will lead to high false negatives
- Similarly, if majority of prefetched domains do not have identifiable keywords, search reconstruction will fail

19

Summary

- Wide-scale study required to fully gauge the effect of DNS prefetching (w.r.t. its privacy implications)
  - Effect on DNS server load remains unclear
- Reduction of user-perceived latency at the cost of privacy
- Primary focus is to foster discussion on the effects of DNS prefetching

20

Questions

21