

# Diversity in DNS Performance Measures

Richard Liston, Sridhar Srinivasan and Ellen Zegura

`liston,sridhar,ewz@cc.gatech.edu`

Georgia Institute of Technology

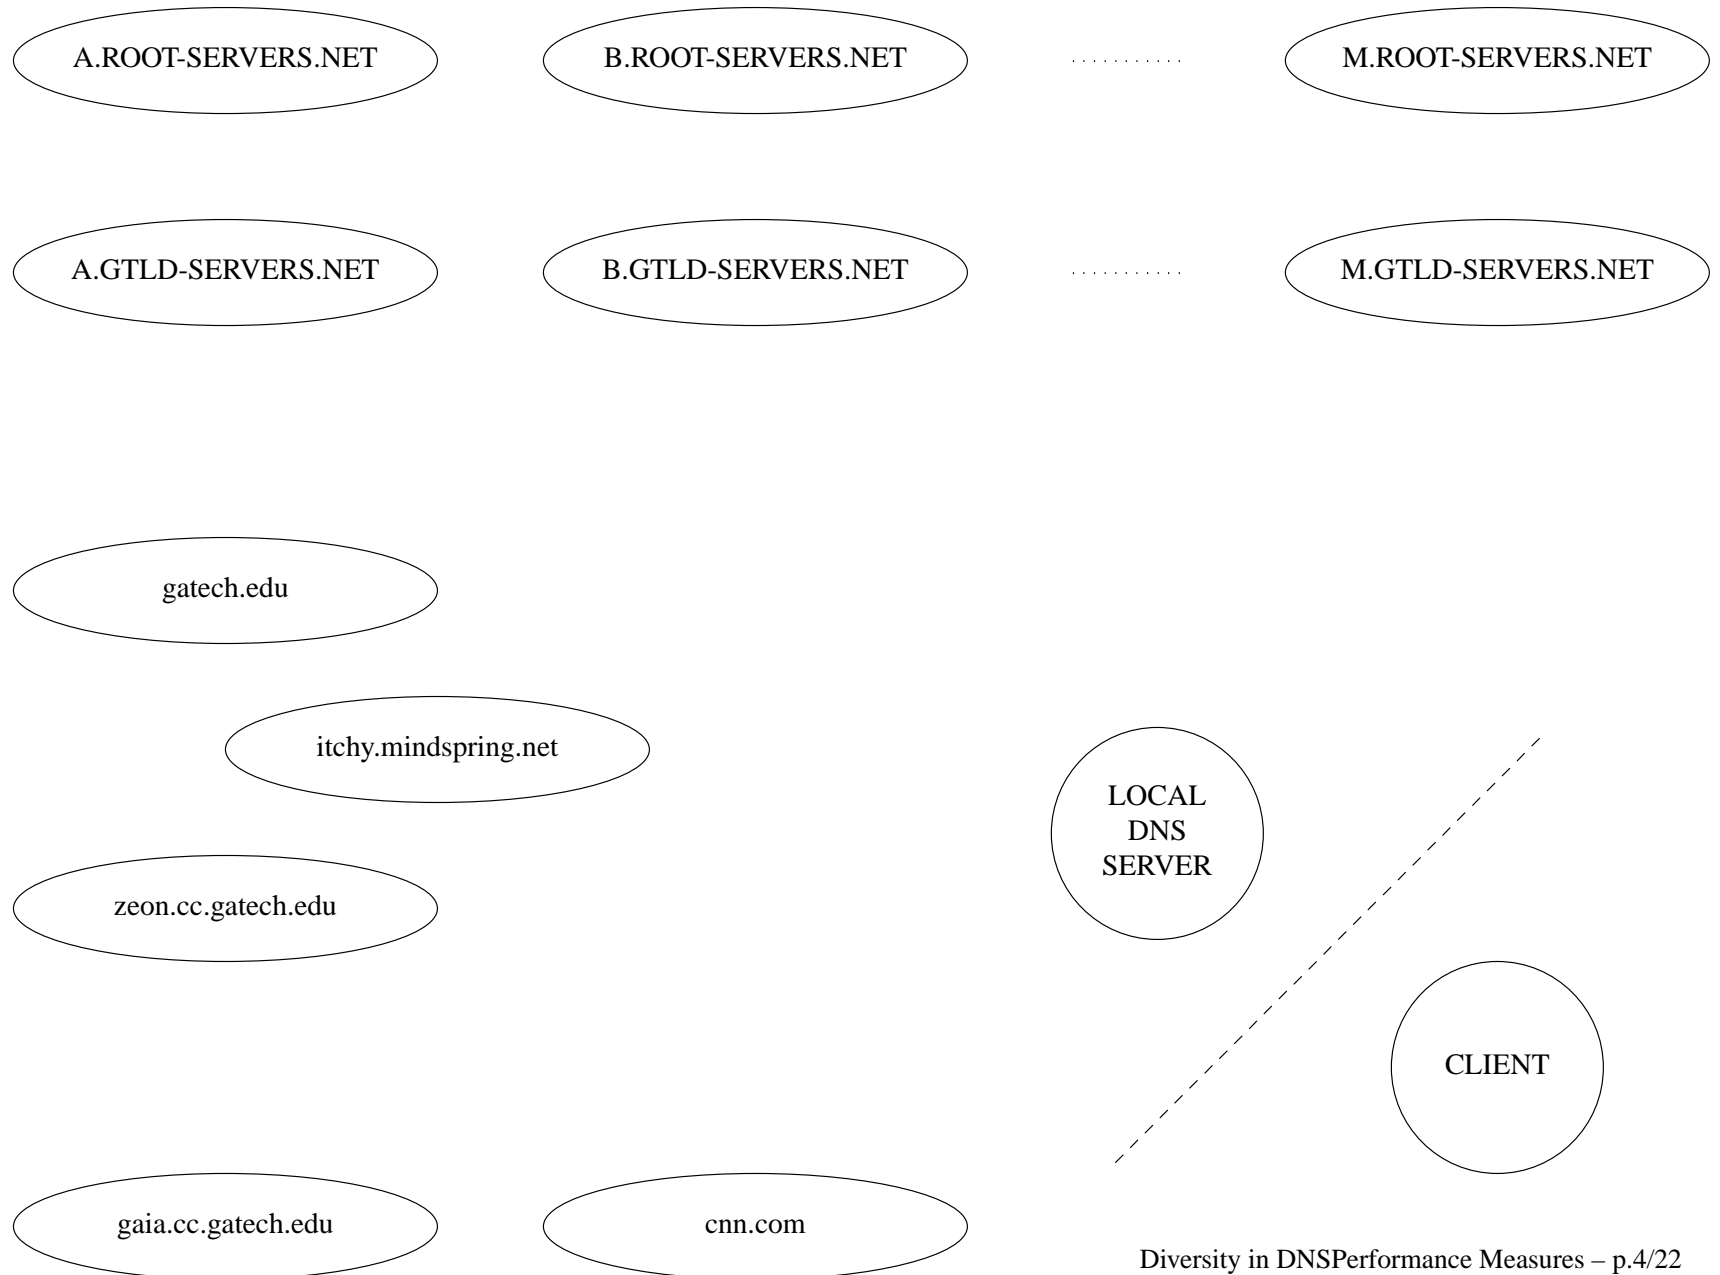
# Roadmap

- Background: DNS
- Problem: Measuring DNS
- Our method: requirements and solution
- Results
- Conclusions

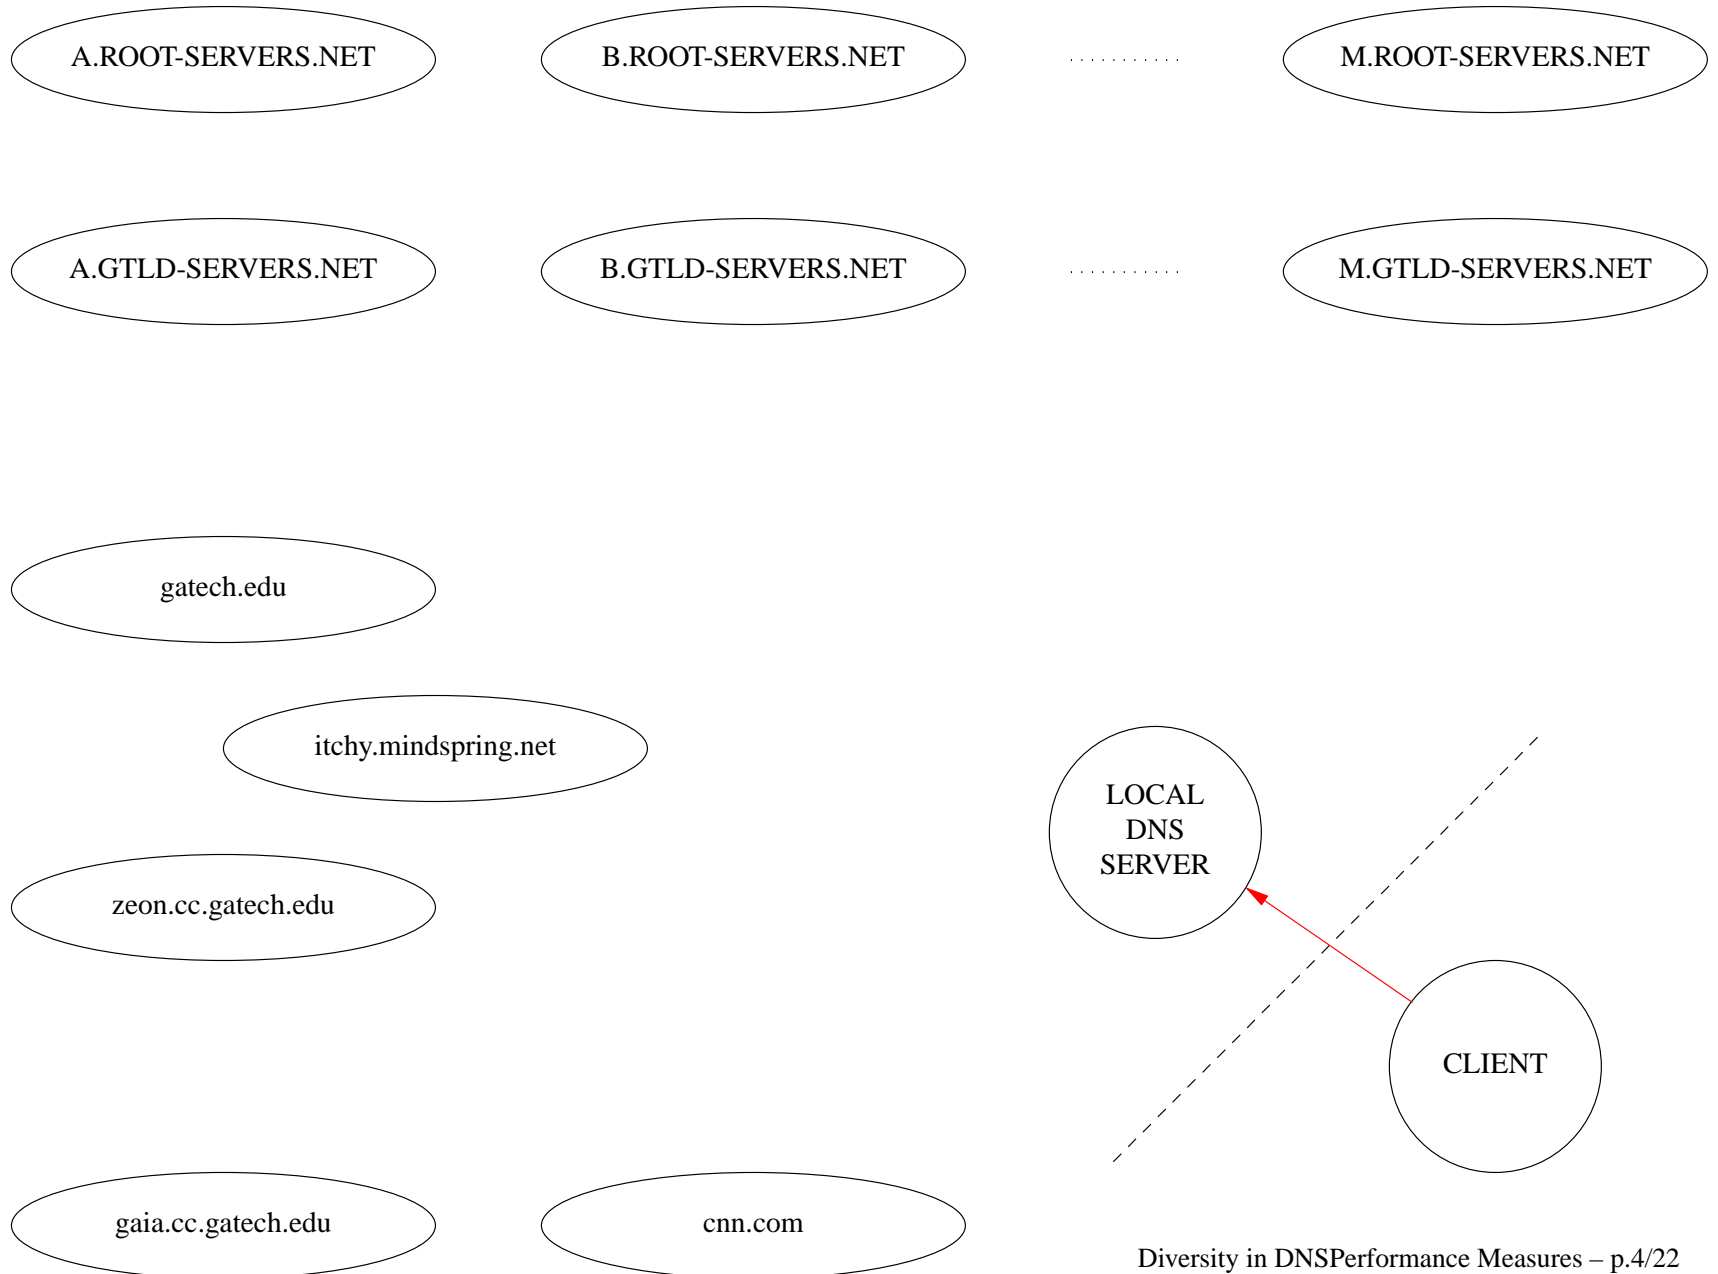
# DNS: Domain Name System

- Solution to the problem of global addressing
- Distributed database over hierarchical namespace
- Allows for local administration of
  - Namespace (via delegation of authority over zones)
  - Caching policies (via TTL)
- Details hidden from users  
(`gethostbyname ( )`)
- Supported by:
  - root servers: .gov, .edu, .fr, ...
  - gTLD servers: .com, .org, ...
  - other servers: lower levels in the namespace hierarchy

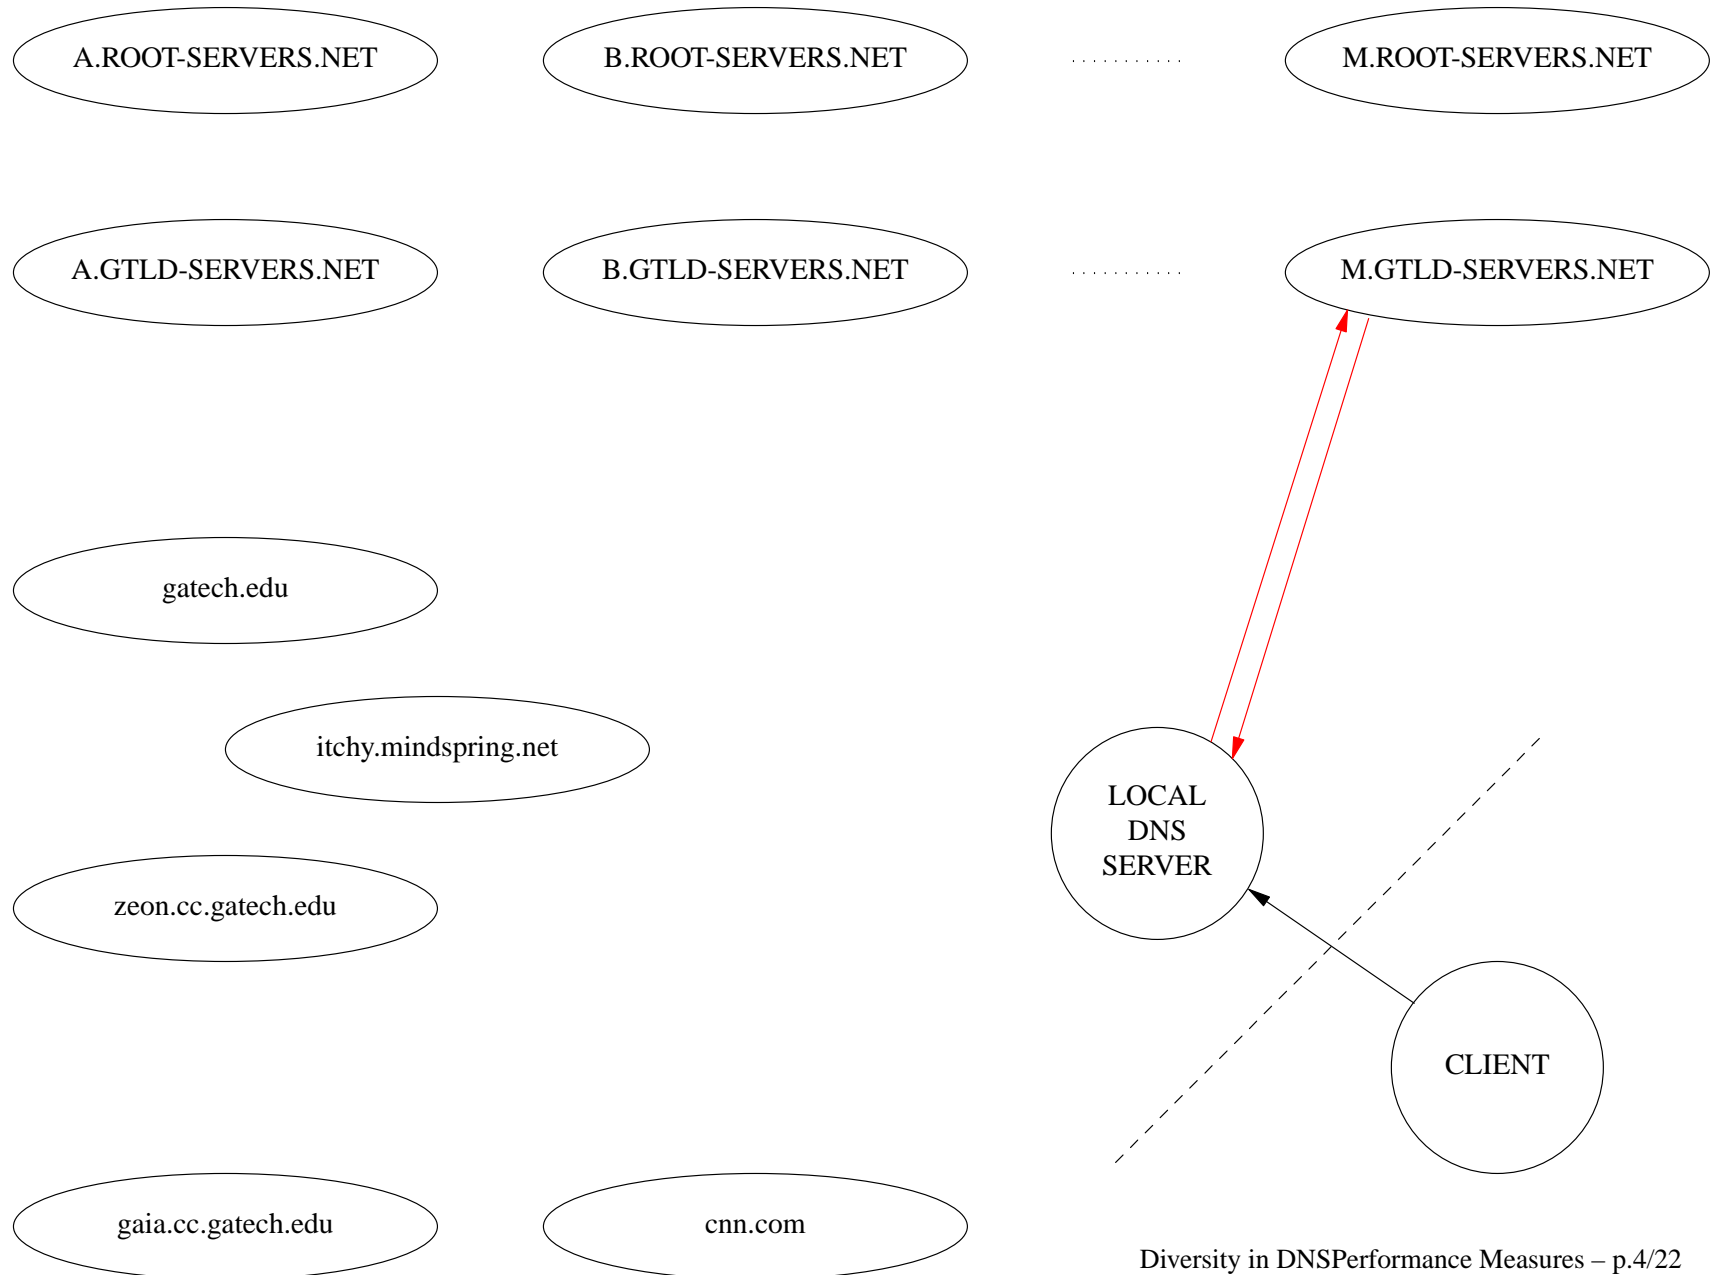
# What Happens in the Wild...



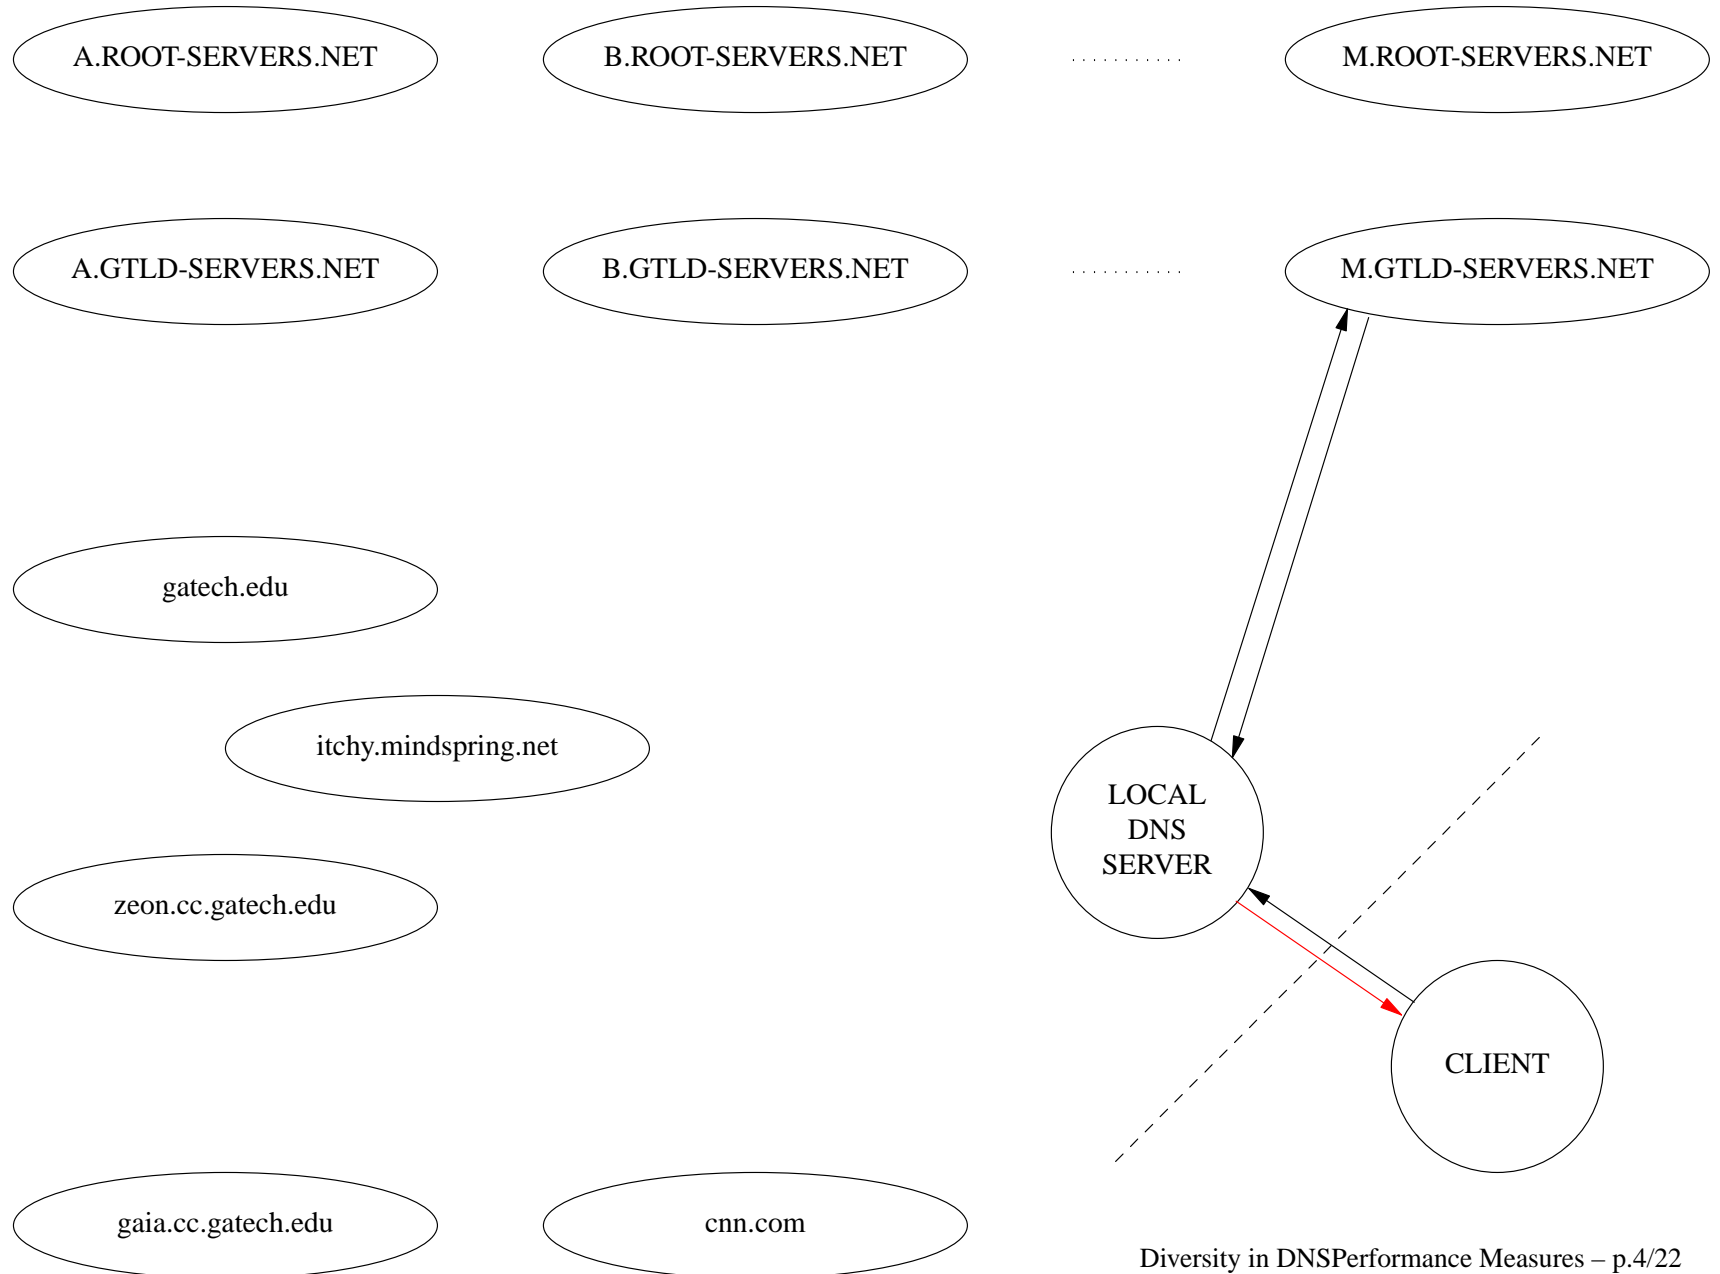
# What Happens in the Wild...



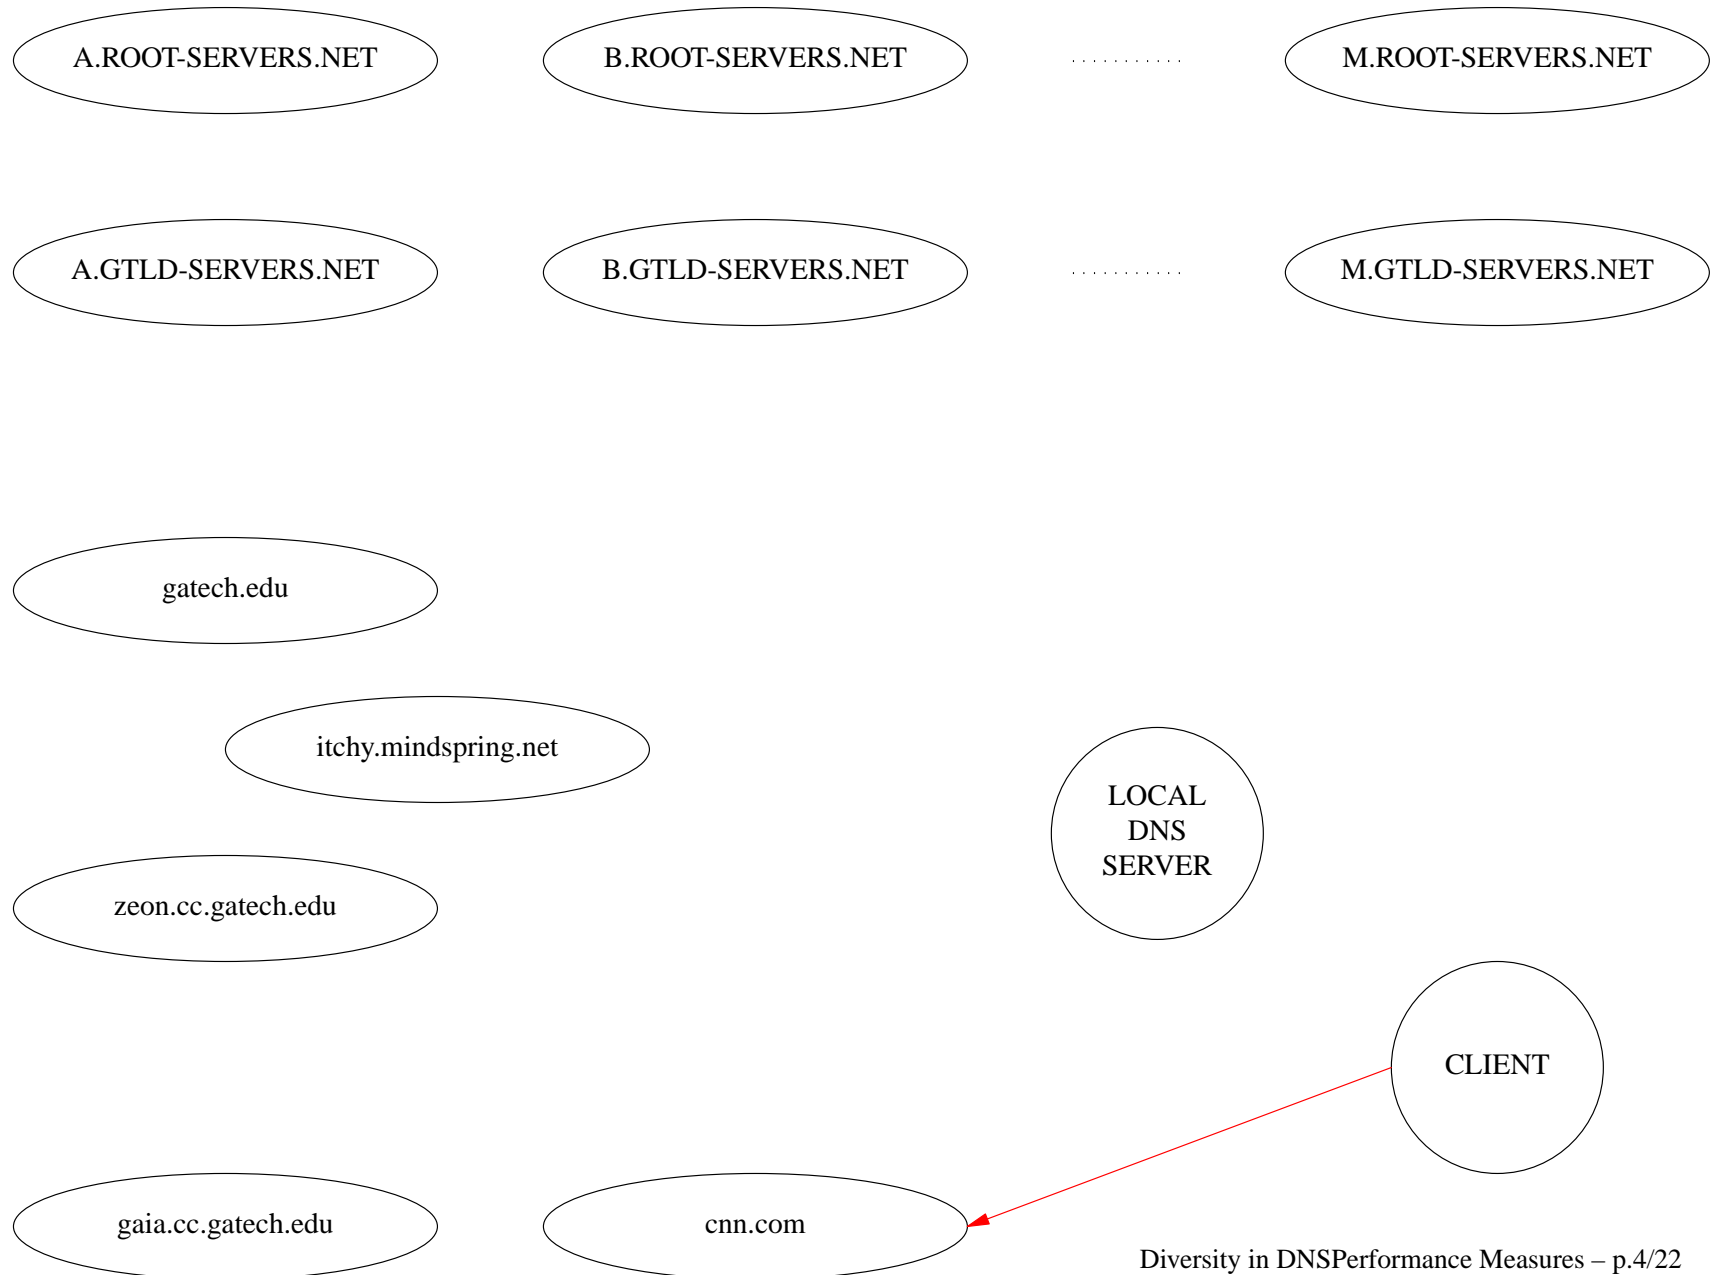
# What Happens in the Wild...



# What Happens in the Wild...




# What Happens in the Wild...

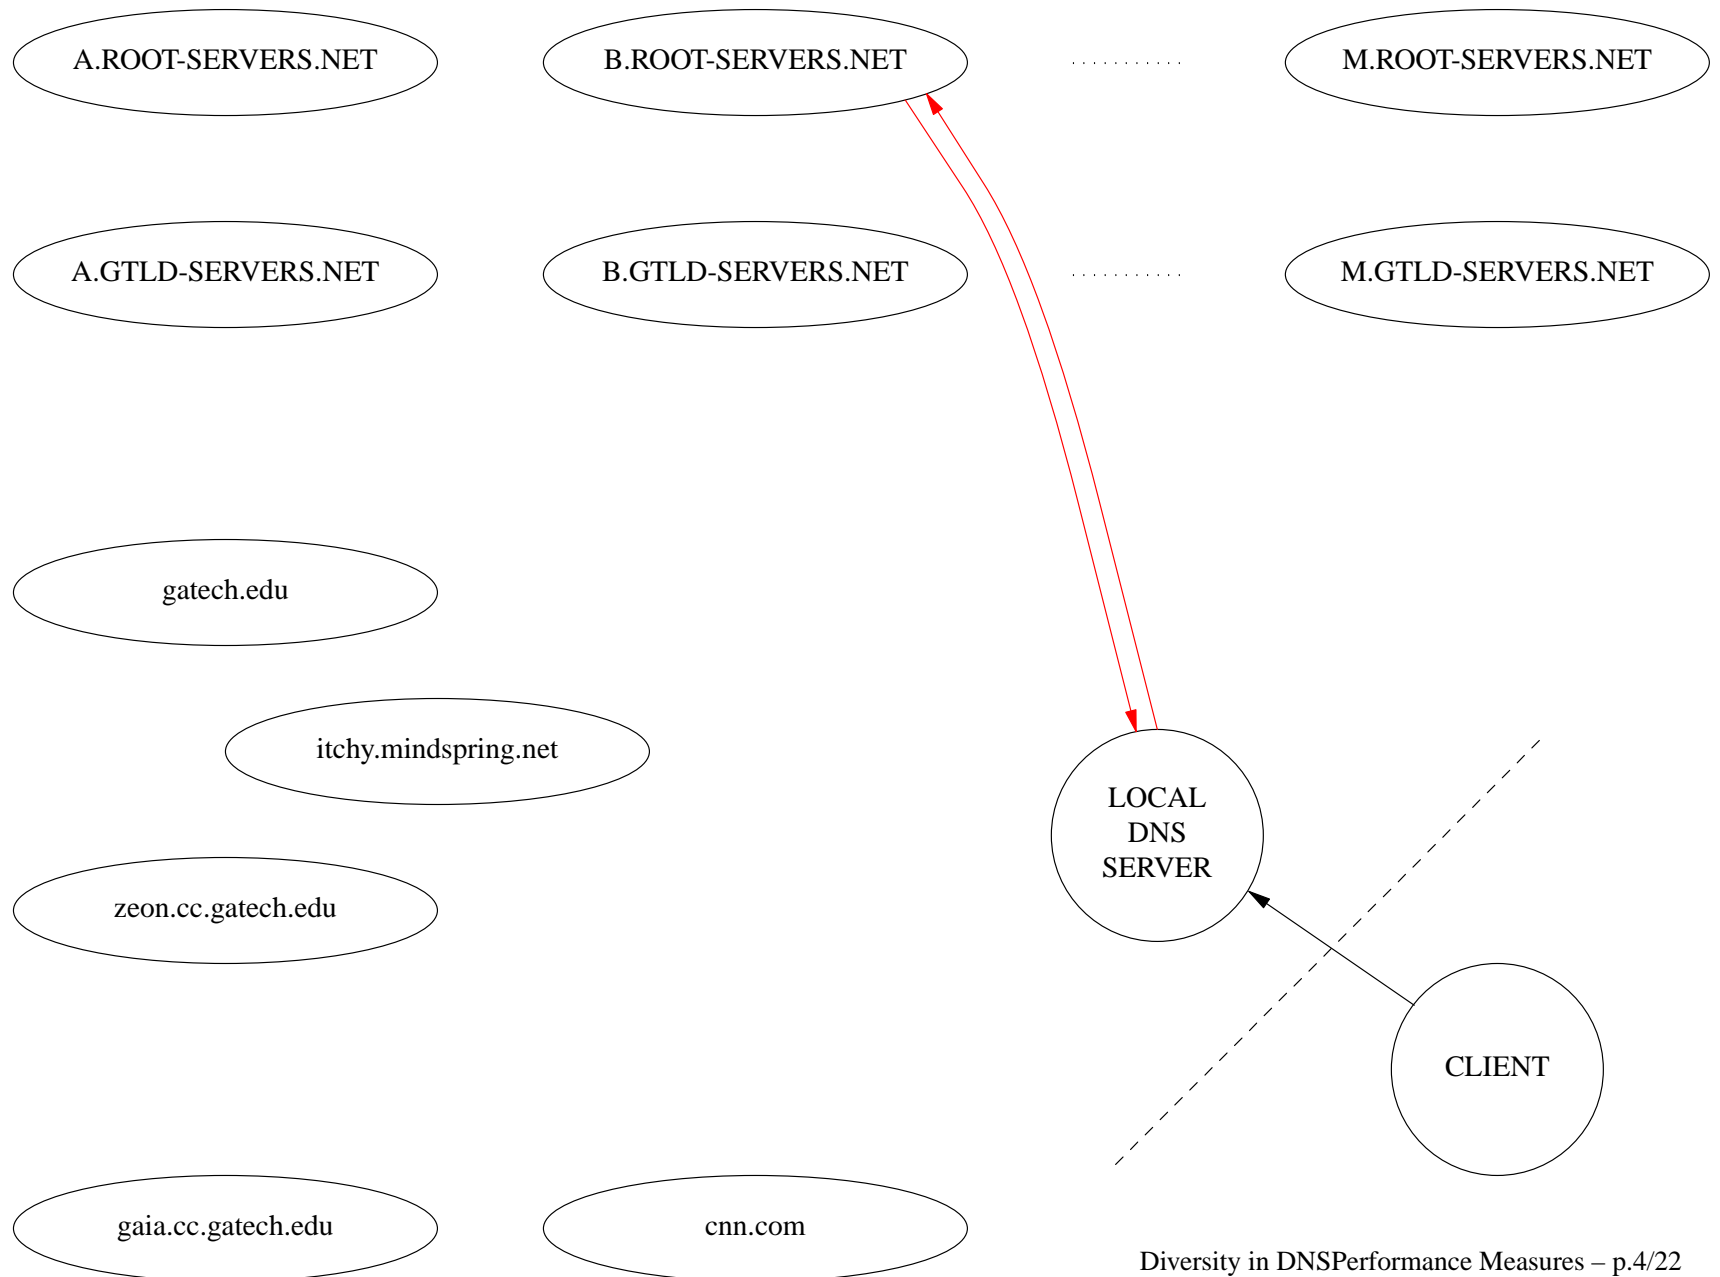




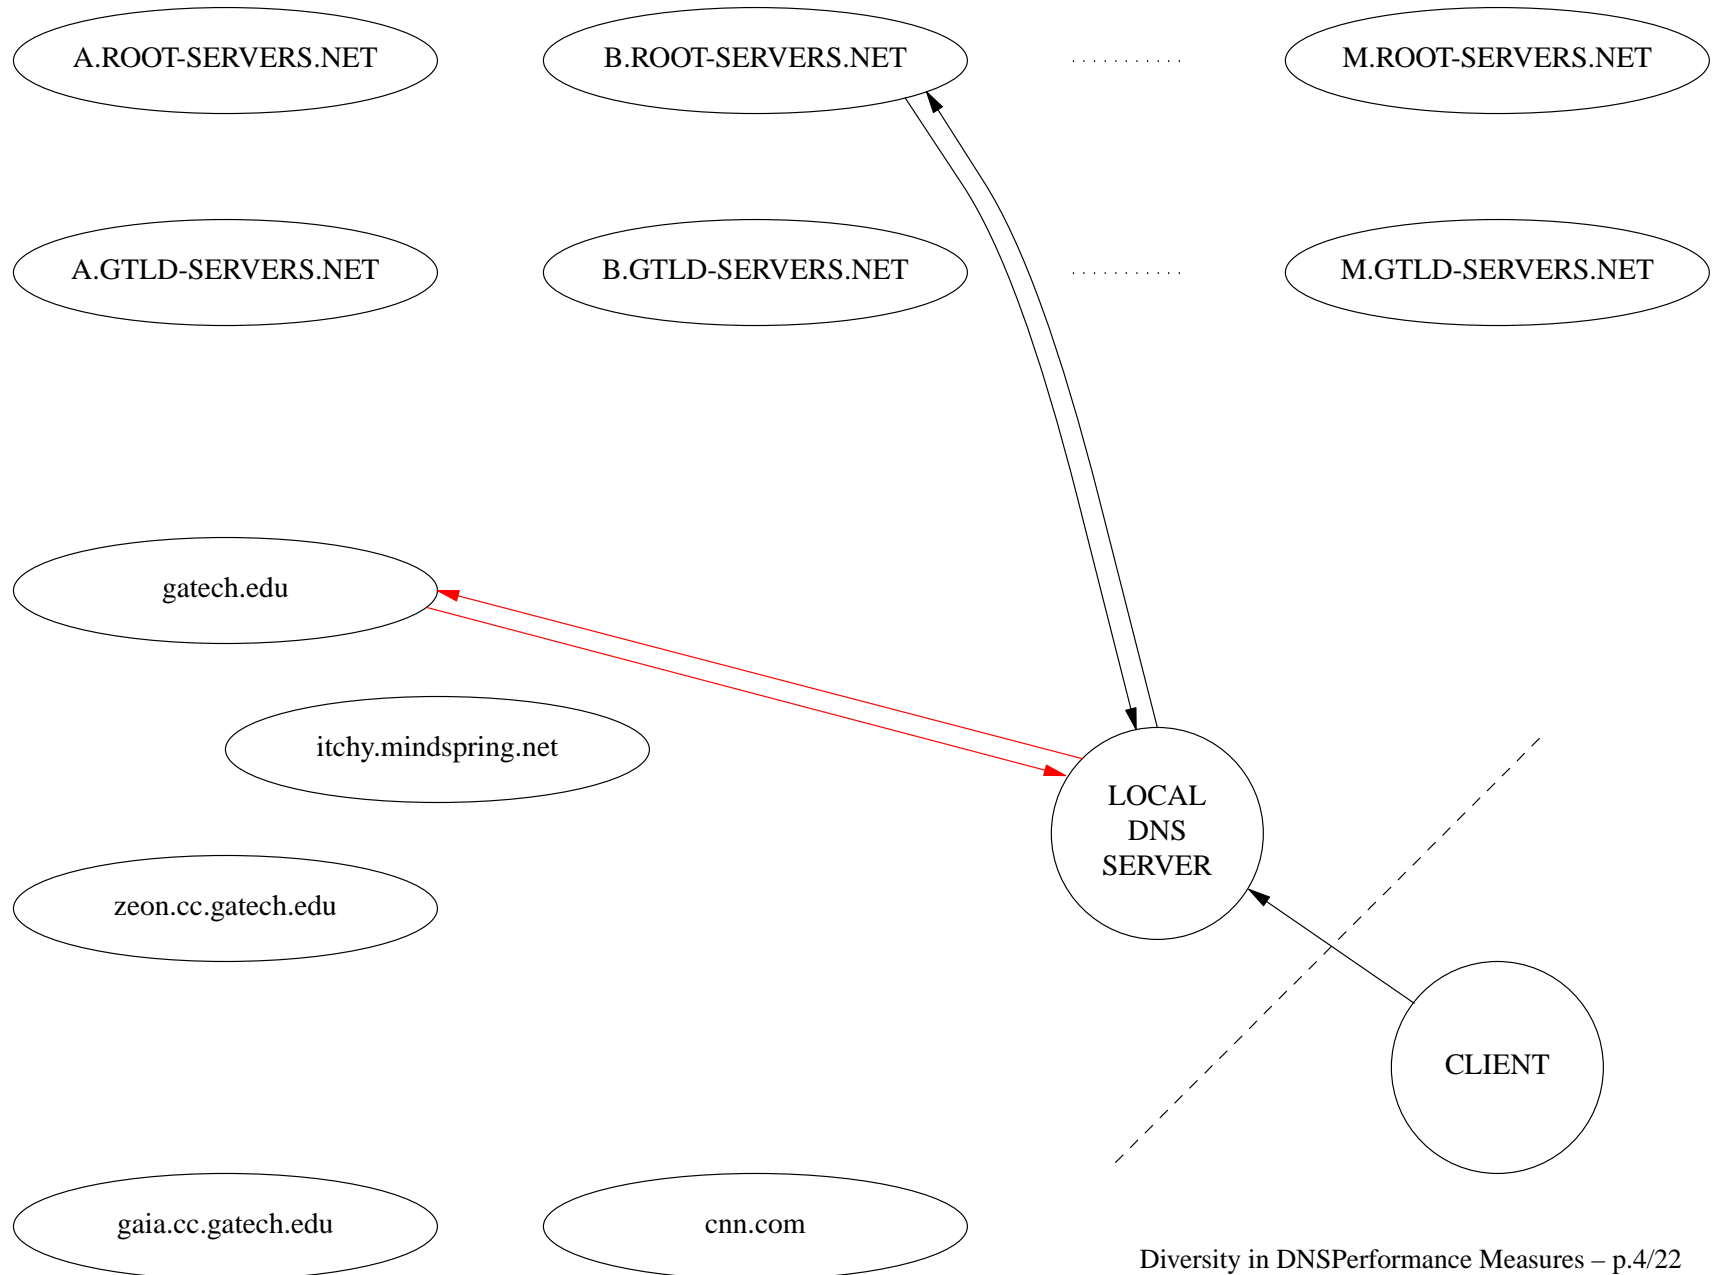
# What Happens in the Wild...



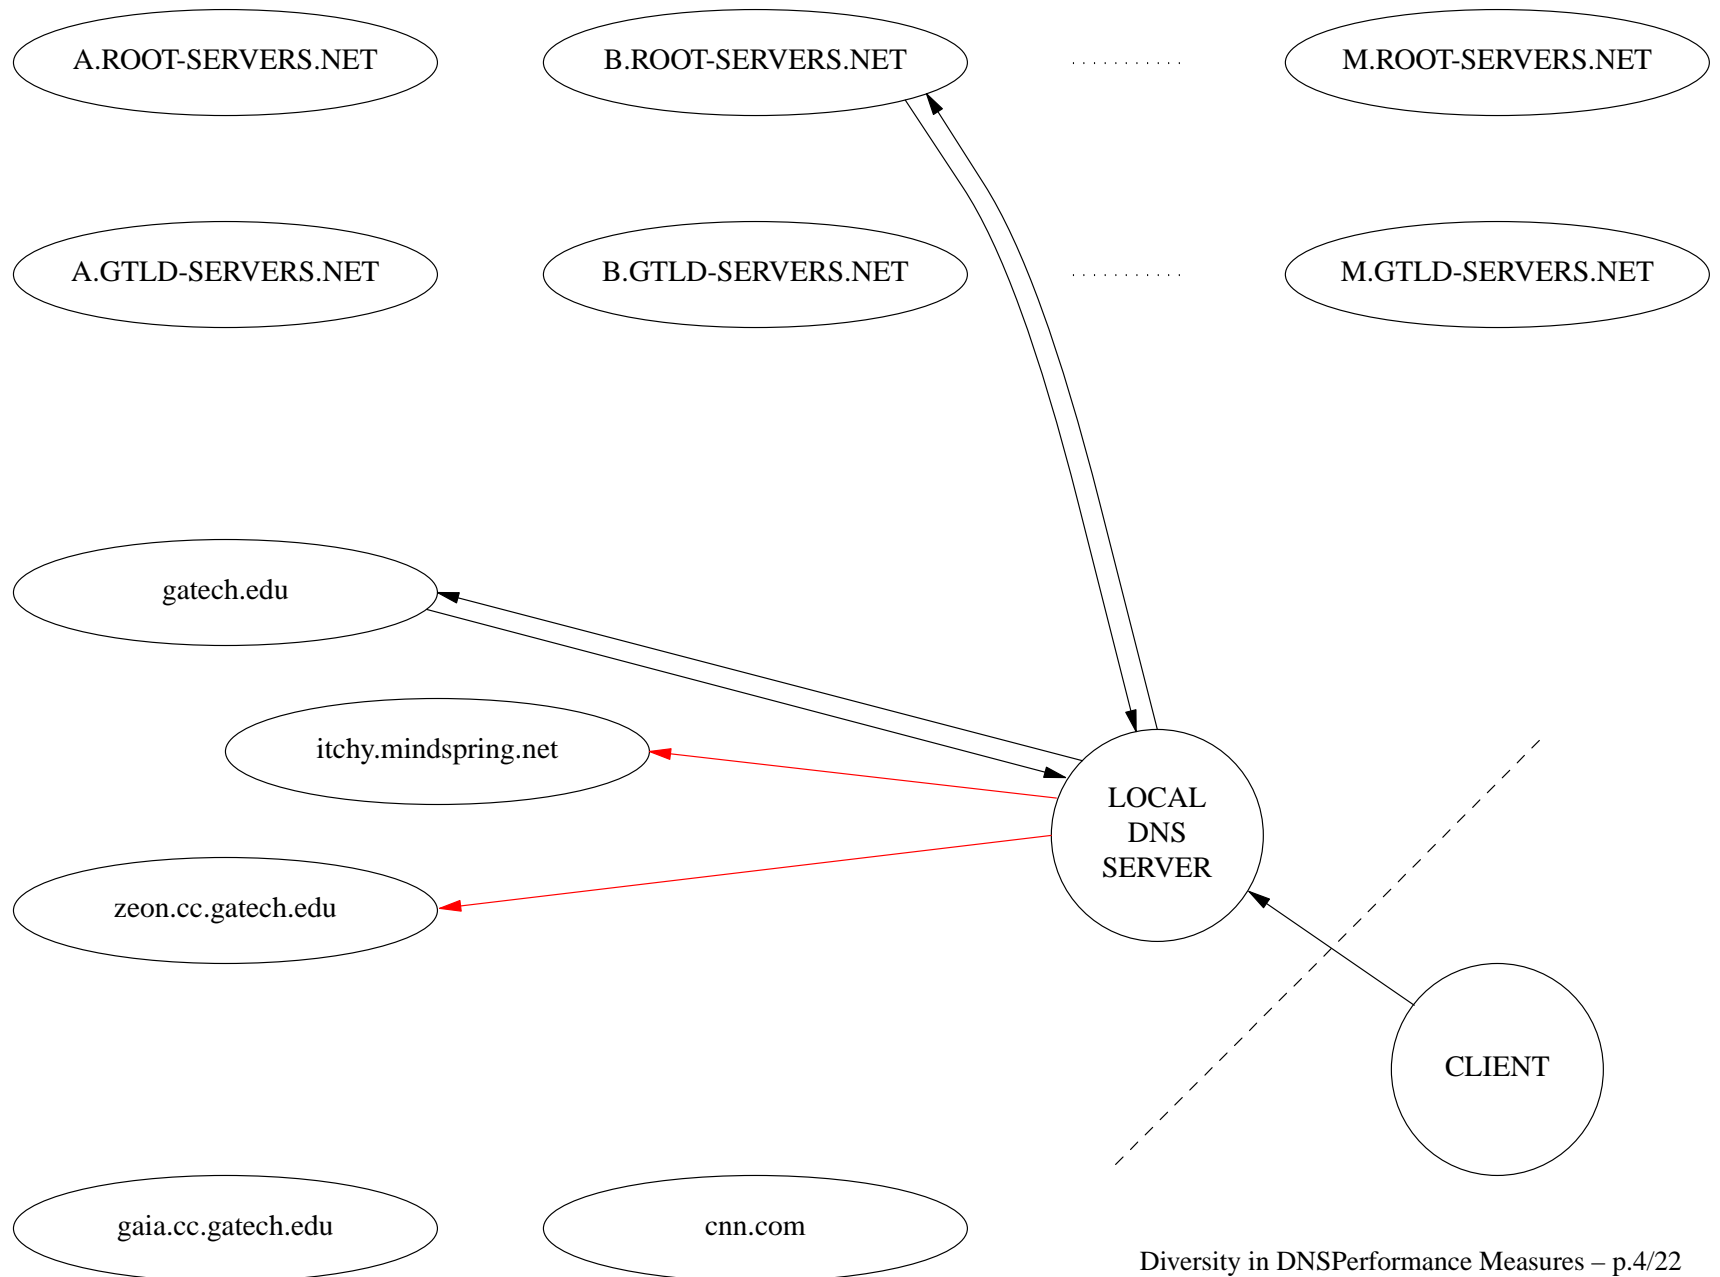
# What Happens in the Wild...



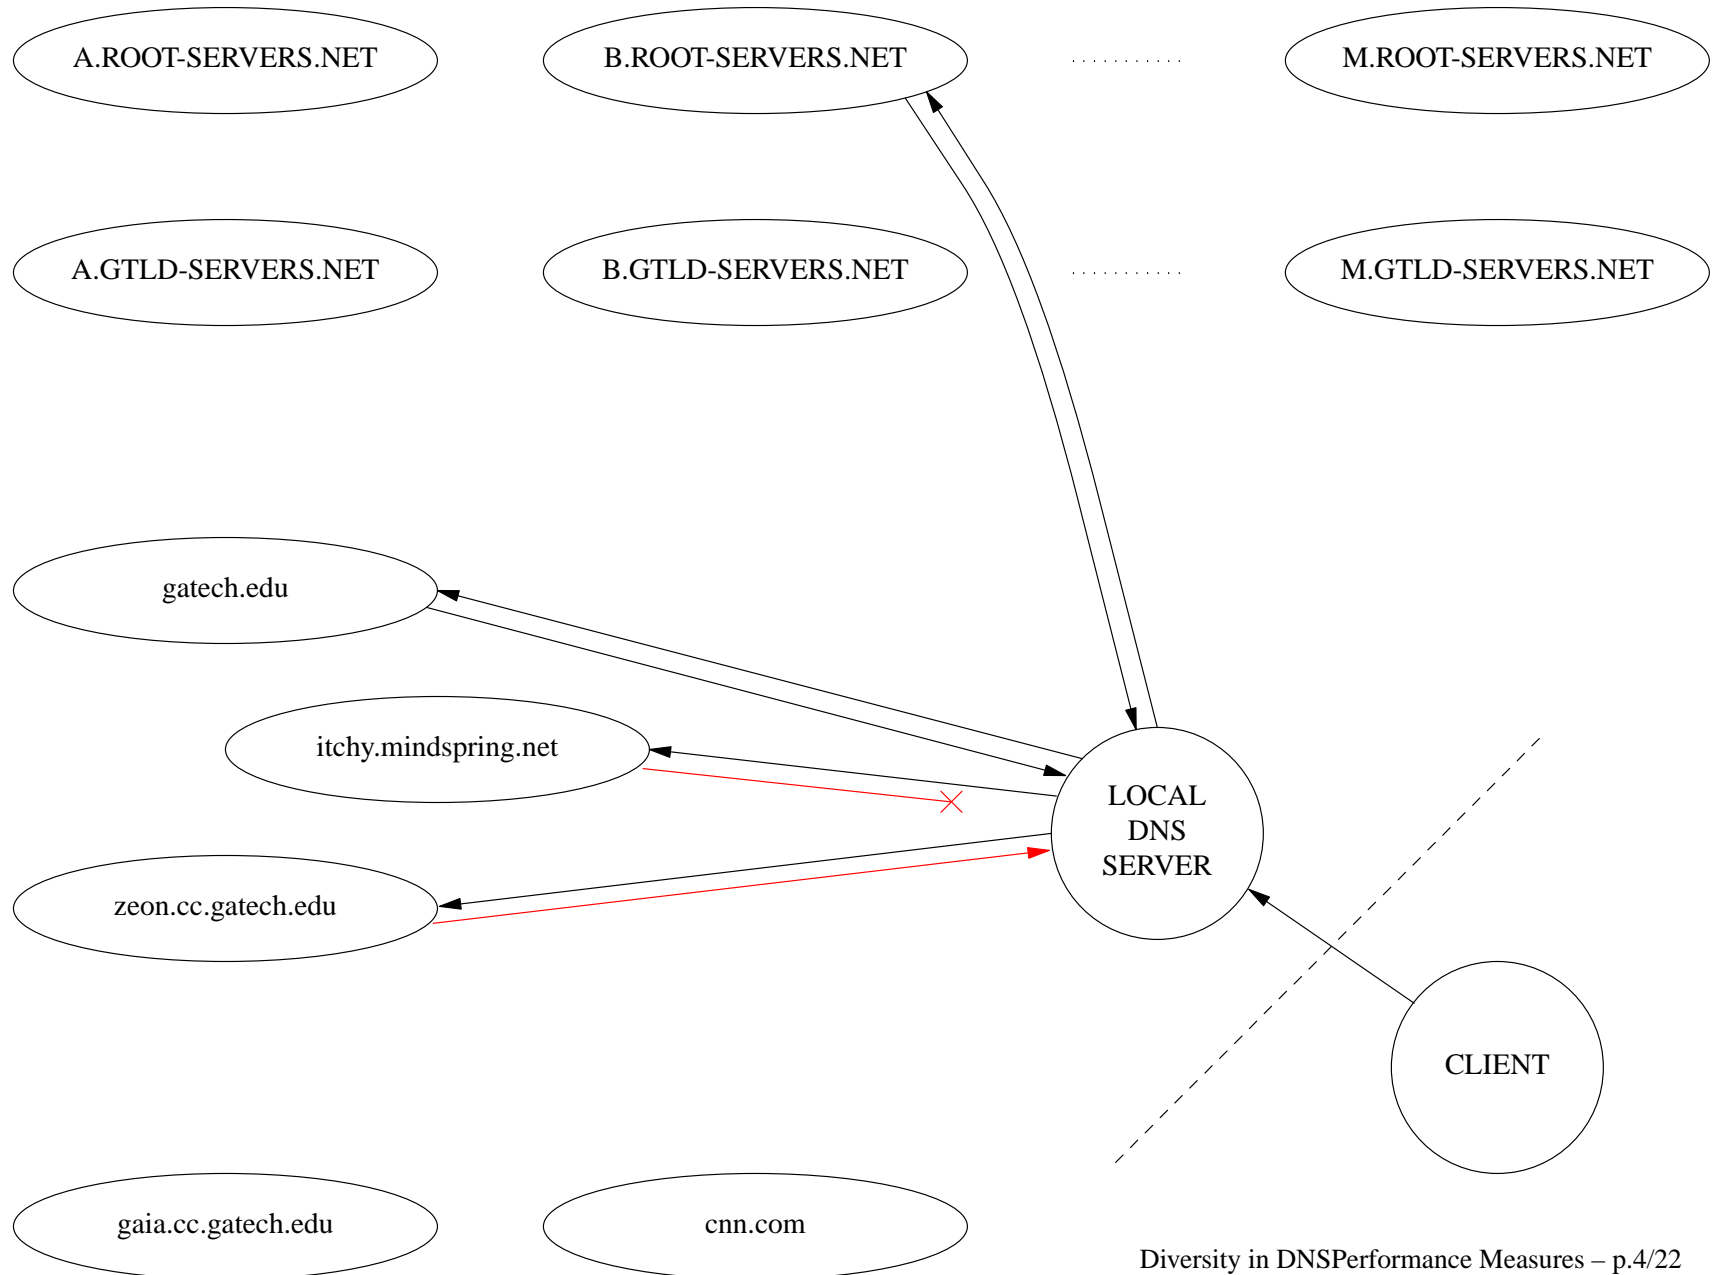
# What Happens in the Wild...



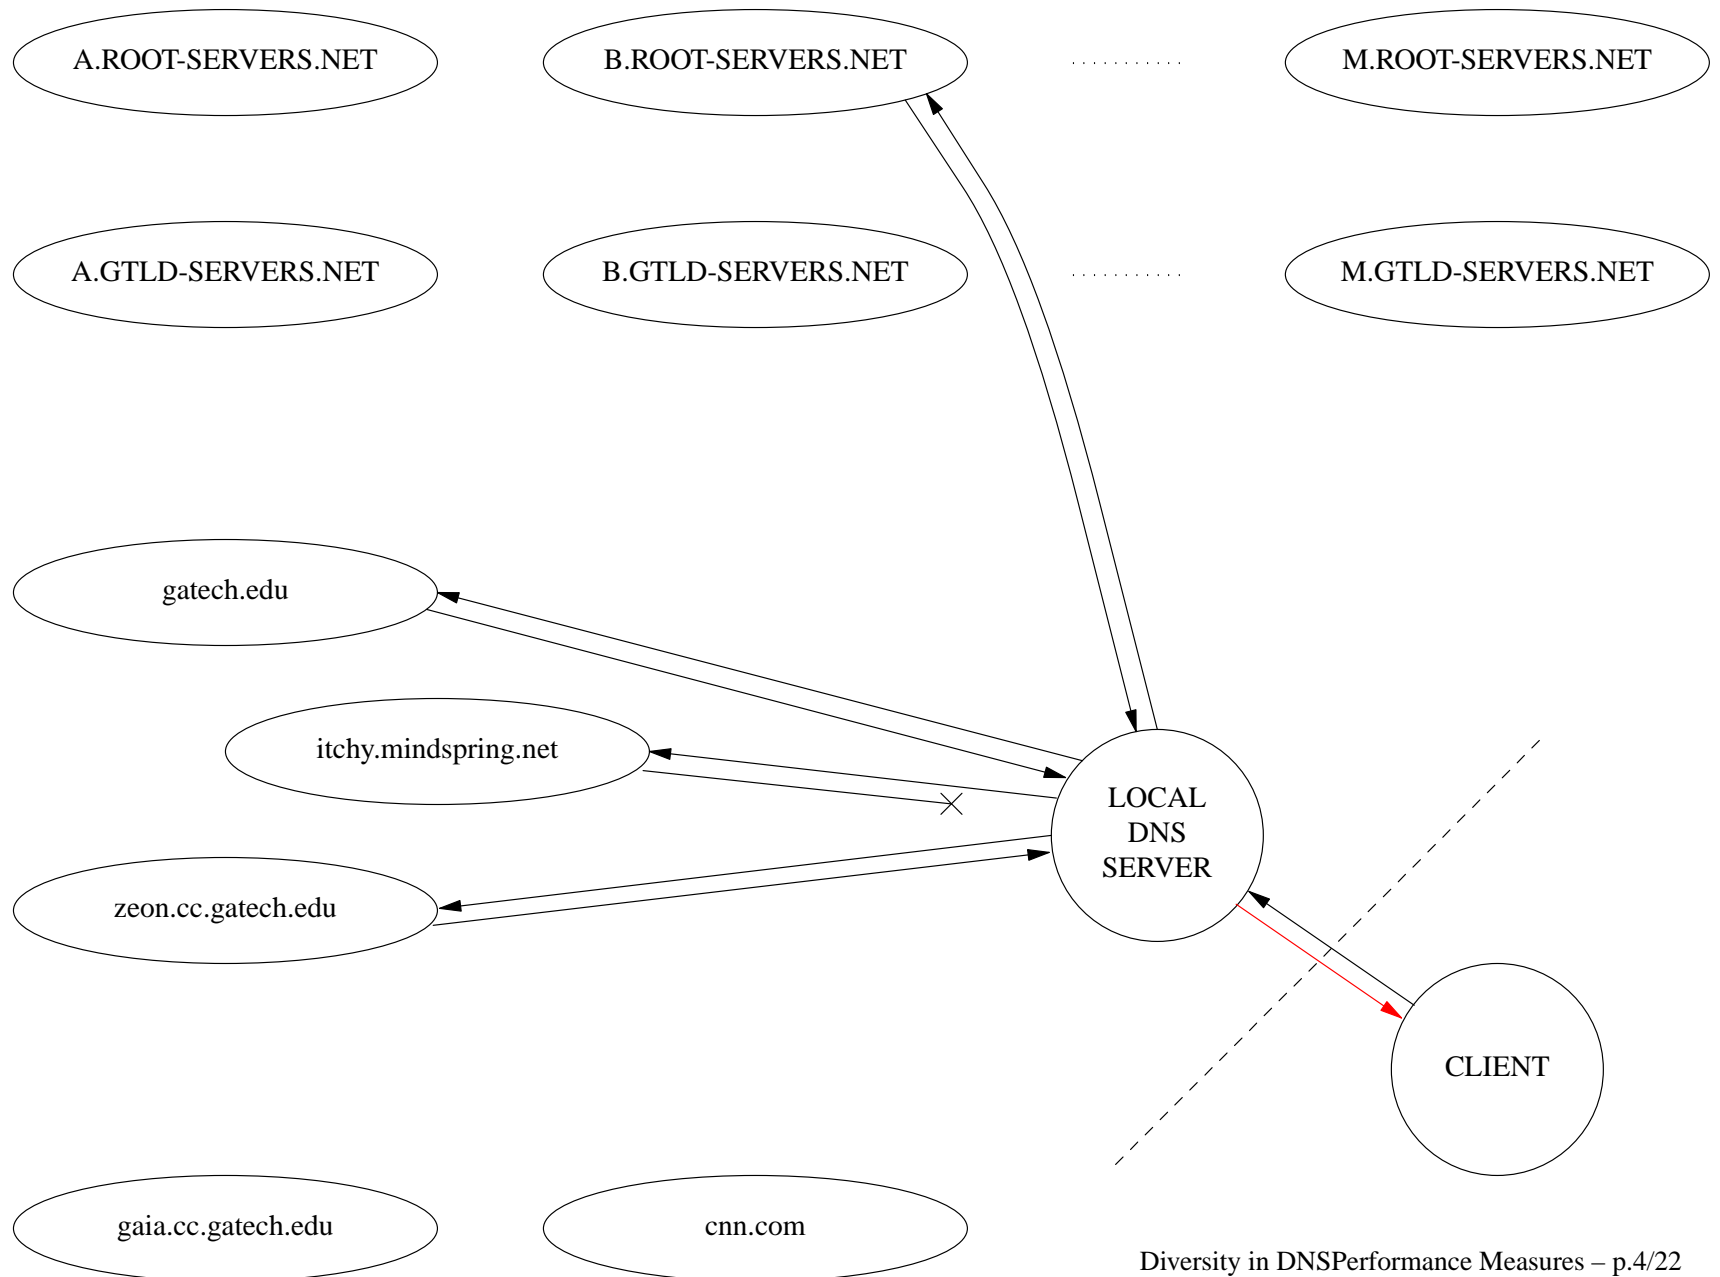
# What Happens in the Wild...



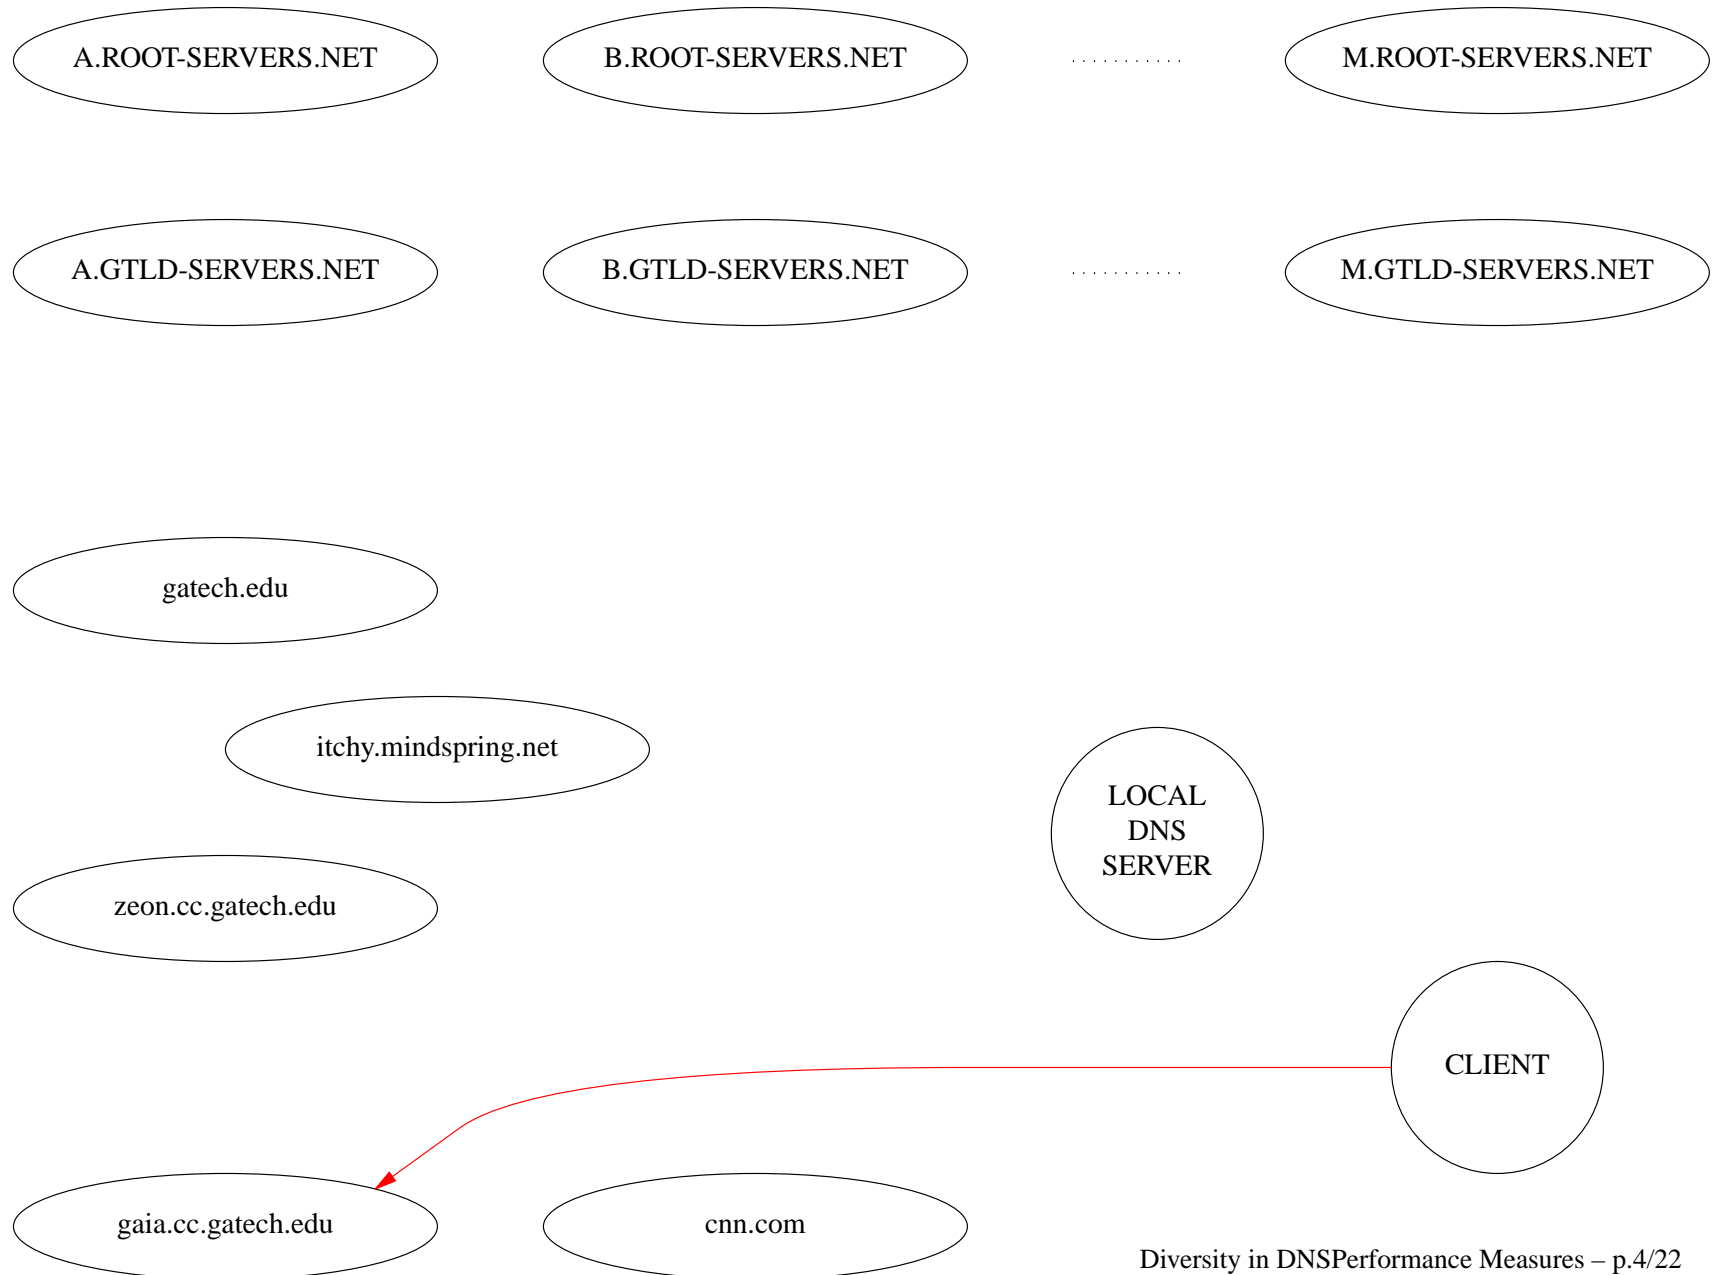
# What Happens in the Wild...



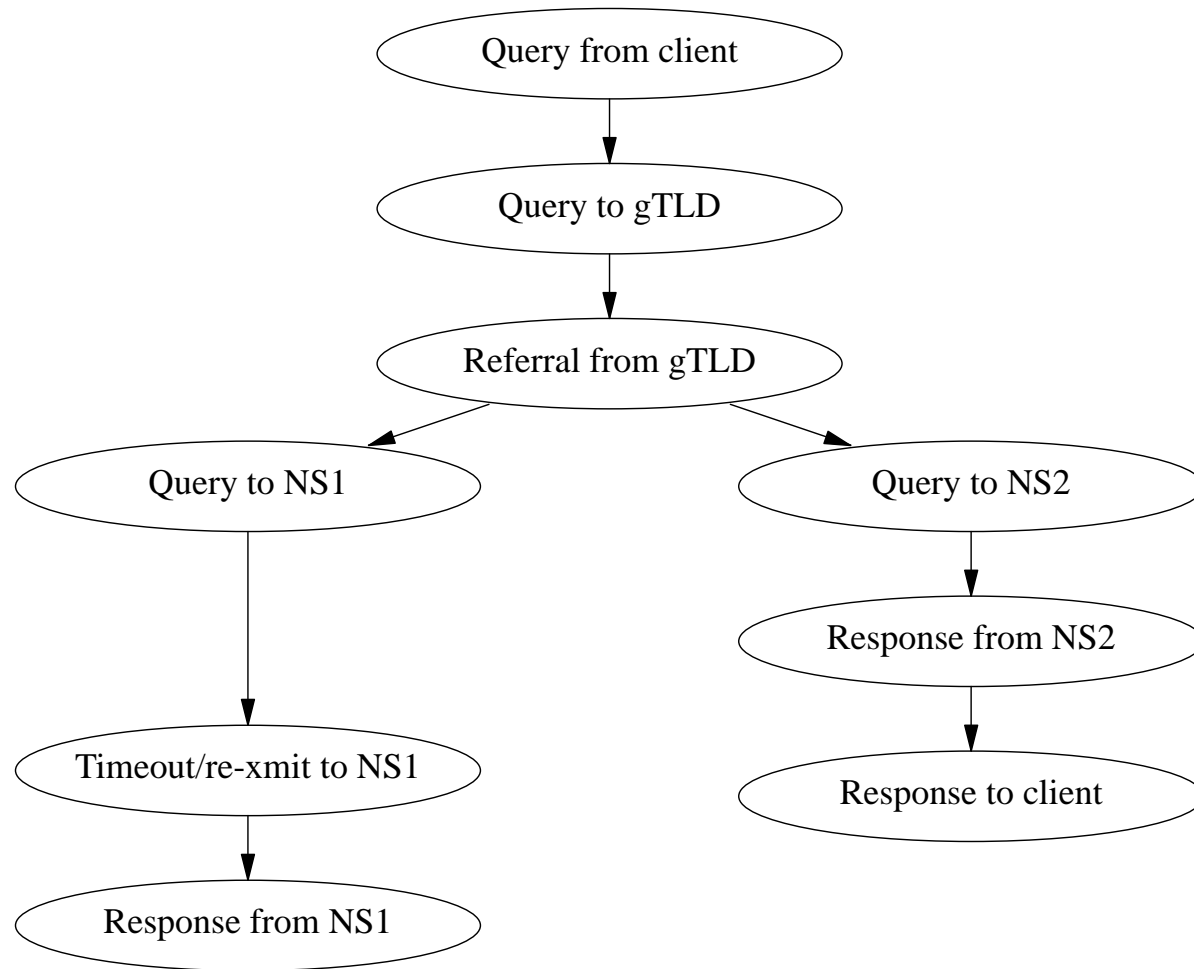
# What Happens in the Wild...



# What Happens in the Wild...

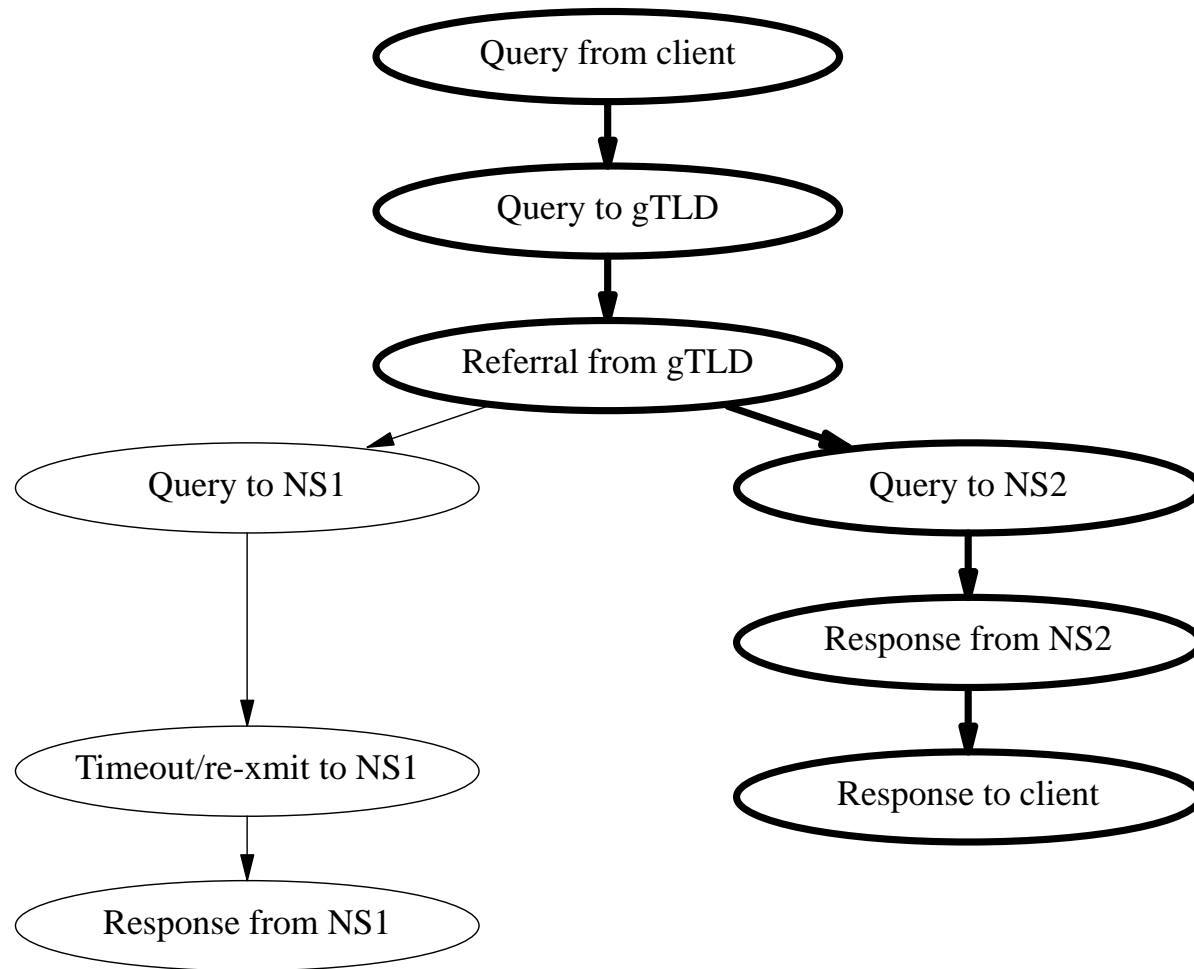


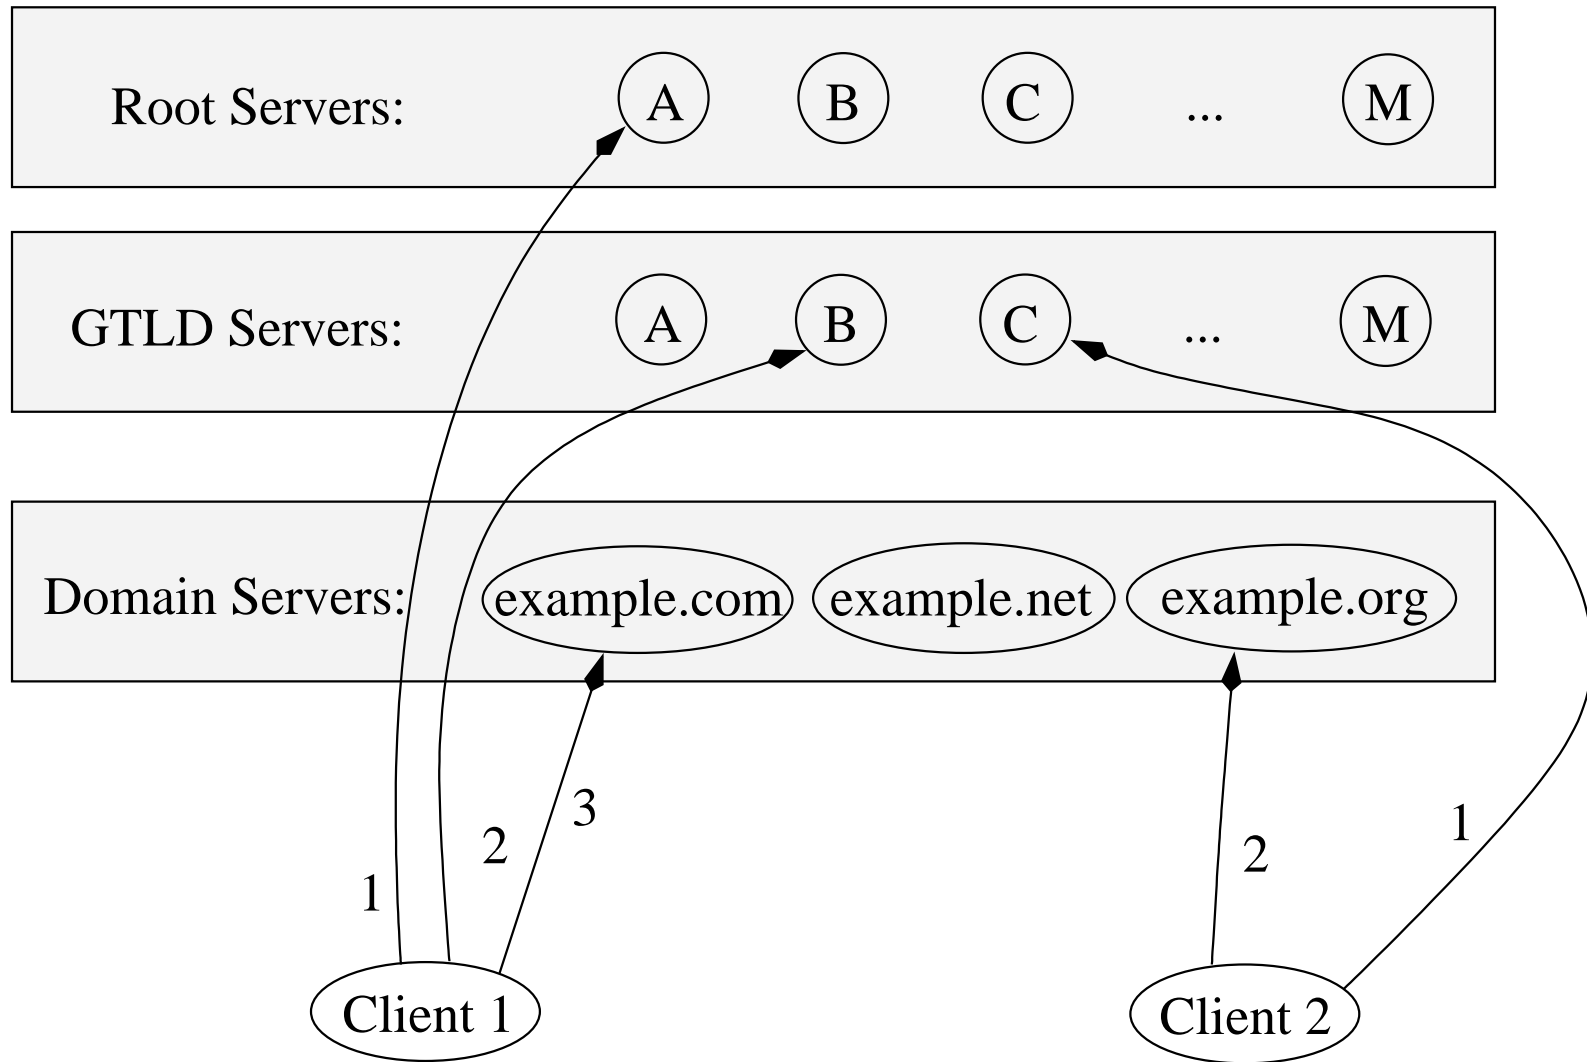
# Resolution Tree: Critical Path

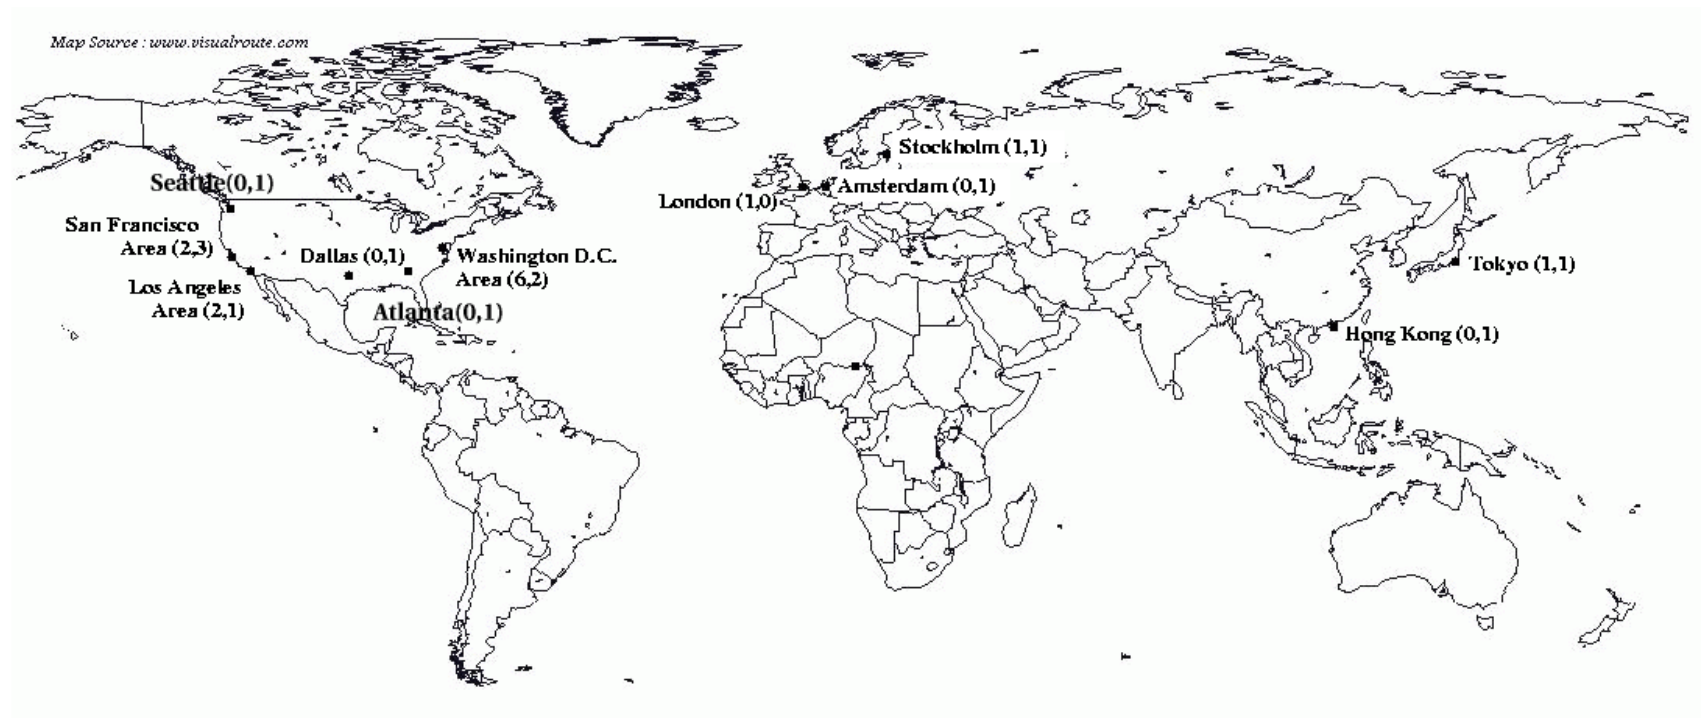




# Resolution Tree: Critical Path







# Measuring DNS

- Primary question in this study: To what extent do DNS performance measures for non-cached names vary across Internet locations?
- Measures:
  - Overall response time
  - Root and gTLD server response time
  - Aliasing via CNAME
  - TTL (Time-To-Live in seconds)

# Measurement Method

- Criteria:
  - Capture fine-grained information about DNS system operation
  - Compare data collected at different sites
  - Easily collect data at multiple locations

# Measurement Method

- Criteria:
  - Capture fine-grained information about DNS system operation
  - Compare data collected at different sites
  - Easily collect data at multiple locations
- Solution:
  - Modify BIND to log relevant info: tx/rx packets, packet type, timeouts, timestamp
  - Package modified BIND w/non-root config files and driving script w/15K domain names
    - Use `dig` to query the modified server
    - Domain names are diverse, non-cached
  - Post-process logs

# Data Collection


- NIMI nodes, colleagues and Linux users; 75 sites in 21 countries and territories (still need China and India!)
- Different days of week and times of day
- Variety of connection technologies: DSL, PPP, cable modem, GigE, etc.
- January 2002 and late March/early April 2002
- Single run: 4 to 6 hours
- Very lightweight

# Validating Workload Set

- Completion and success rates
- Response codes:
  - 0: No error
  - 2: Remote server failure
  - 3: Name does not exist
- Completed (0 or 3): 96.4% to 98.1%
- Successful (0): 92.7% to 94.7%



# Mean Response Time, completed (MRTc)



Range from 0.95 seconds to 2.31 seconds!

# Measuring Major Factors

- Connectivity
  - MINc: Minimum response time for completed lookups
  - Minimum round trip time to closest name server queried

# Measuring Major Factors

- Connectivity
  - MINc: Minimum response time for completed lookups
  - Minimum round trip time to closest name server queried
- Loss Rate
  - Total retries along all critical paths

# Measuring Major Factors

- Connectivity
  - MINc: Minimum response time for completed lookups
  - Minimum round trip time to closest name server queried
- Loss Rate
  - Total retries along all critical paths
- Root/GTLD server performance
  - MRTr/MRTg: mean response time to root/gTLD servers


# Measuring Major Factors

- Connectivity
  - MINc: Minimum response time for completed lookups
  - Minimum round trip time to closest name server queried
- Loss Rate
  - Total retries along all critical paths
- Root/GTLD server performance
  - MRTr/MRTg: mean response time to root/gTLD servers
- Location relative to other nameservers
  - MRTl: mean response time to a fixed set of last nameservers contacted

# Correlations

Factor	Correlation w/MRTc
Connectivity (MINc)	0.62
Loss Rate (# retries)	0.50
Root Server Performance (MRTr)	0.86
GTLD Server Performance (MRTg)	0.94
Location (MRTl)	0.90

# Impact of Root/GTLD/Other



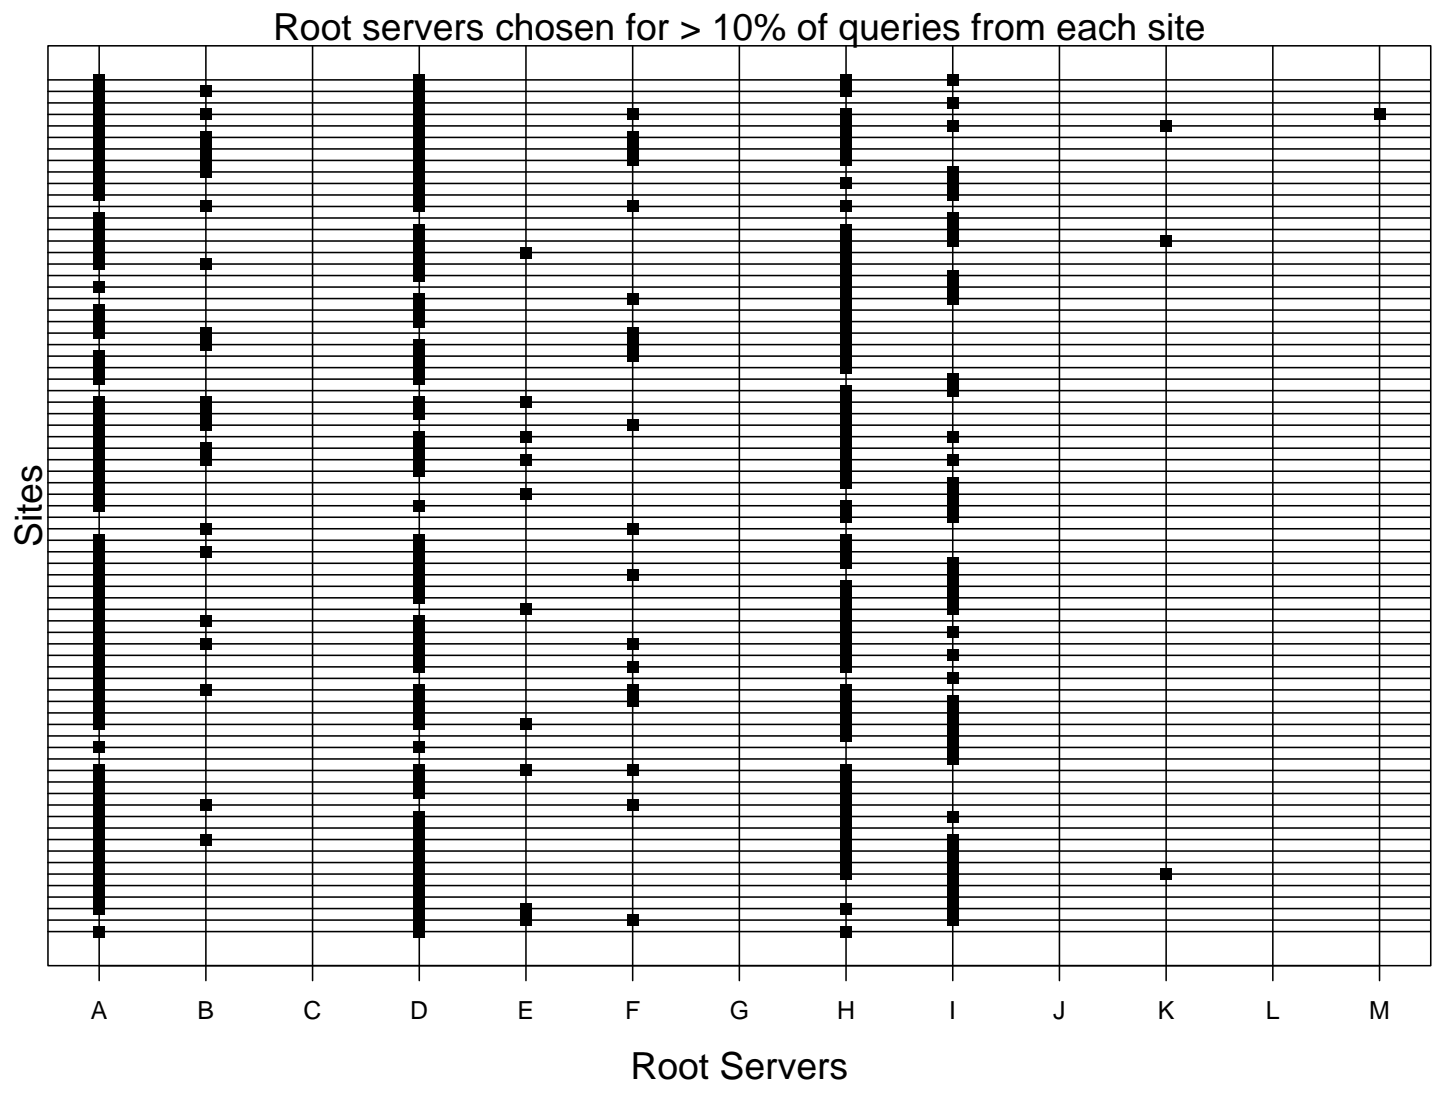
# Some Implications

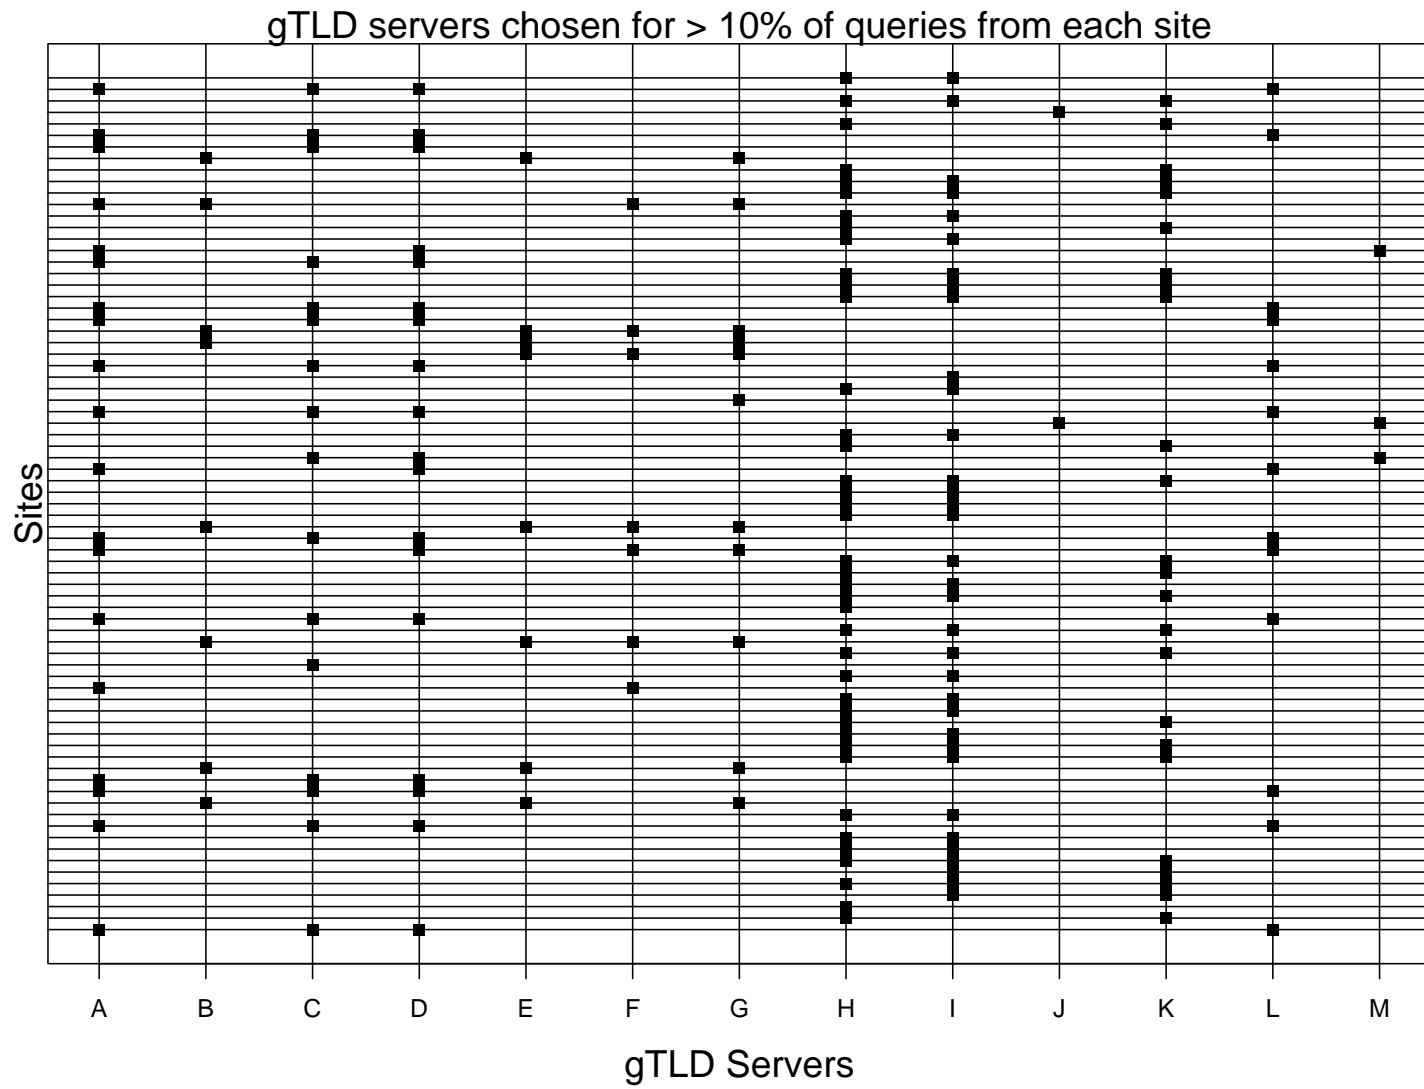
- Roots: poor performance barely noticeable
- GTLDs: poor performance more noticeable (~20% - 30% of time spent in querying gTLD servers); better server placement could help
- Other servers: most potential for performance enhancement; do CDNs want to provide DNS-only services???
- A service differentiator for ISPs may involve performance from gTLD servers
- ISPs may consider choosing different routes to gTLD servers



# More on root/gTLD Interaction

- BIND employs a server selection algorithm
- Maintains history of NS response times
- Ages performance history to sample all servers
- Converges on one or more servers
- "Favored" servers:  $\geq 10\%$  of queries





# Other Results: CNAME/TTL

Measure	Observations
CNAME	<ul style="list-style-type: none"><li>• Fairly consistent use across sites</li><li>• 26% of our names were aliases</li><li>• 6% resolved to more than one different CNAME</li><li>• &lt;1% (four) resolved to more than ten different CNAMEs</li></ul>
TTL	<ul style="list-style-type: none"><li>• Distribution the same across all sites</li><li>• Most popular TTLs: one day, one hour, two hours, 12 hours, two days</li></ul>

# Conclusions

- Most widespread measurement of DNS performance to date
- Measures controlled by admins show low variation: CNAMEs, TTLs
- Measures subject to network location vary widely
  - Mean response time for completed lookups
  - Mean response time from root/gTLD servers
  - Favored root/gTLD servers
- Potential for performance enhancement: gTLD/other servers
- Demonstrated need for multiple measurements

- Still collecting data!
- <http://www.cc.gatech.edu/~liston/dnsperf.html>

Questions?