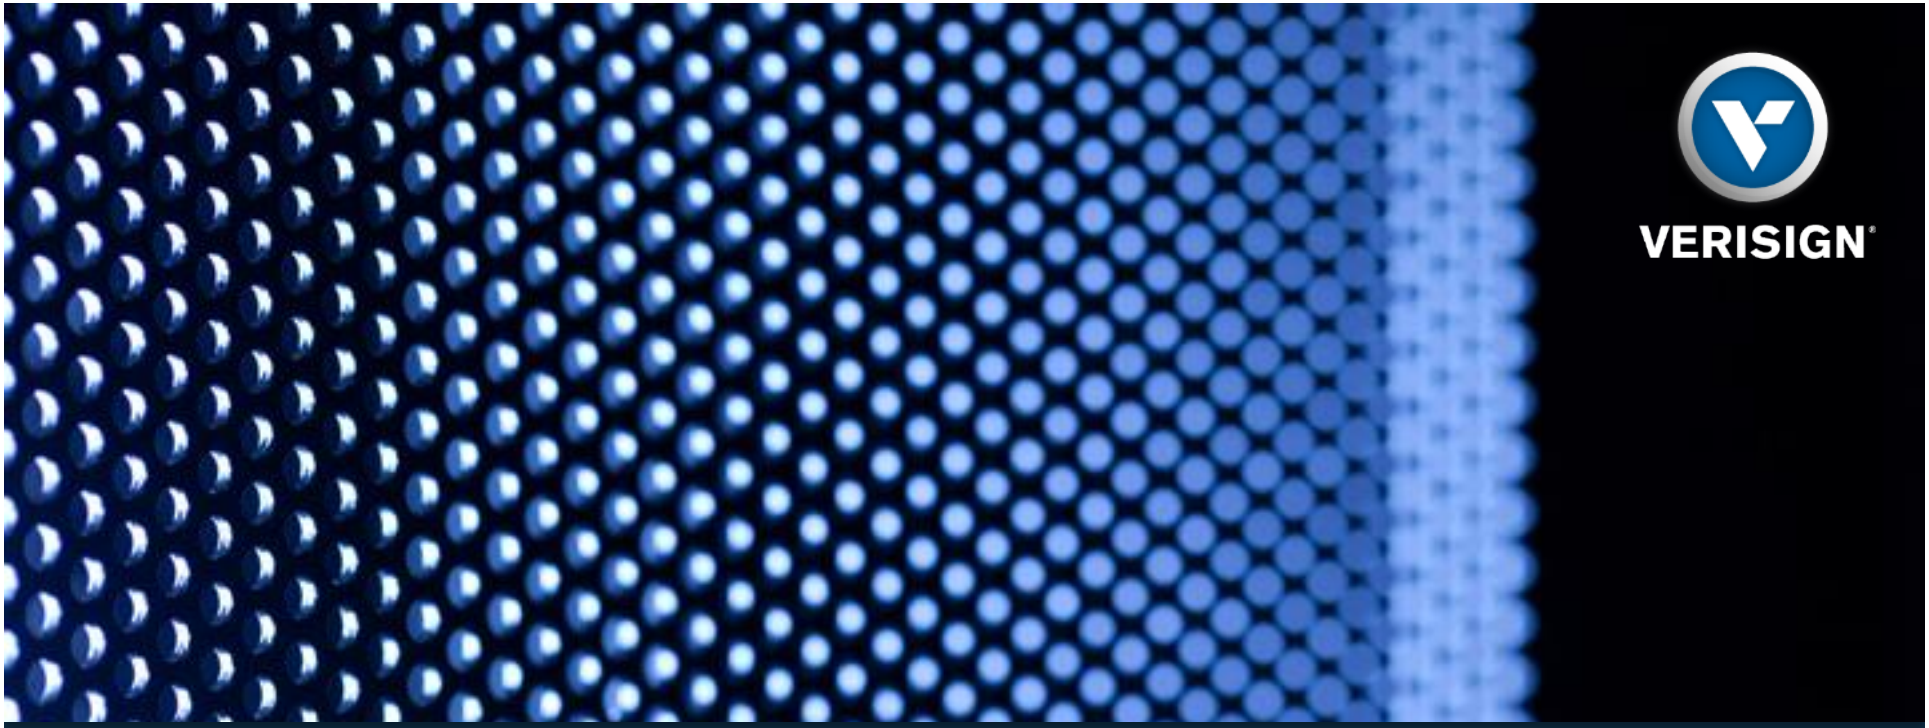

A Multi-perspective View of DNS Availability and Resilience

Casey Deccio, Verisign Labs

AIMS 2015

April 2, 2015

DNS Background/Architecture

DNS Name Resolution

- ***Resolvers*** query ***authoritative servers***
- Queries begin at root zone, resolvers follow downward referrals
- Resolver stops when it receives authoritative answer

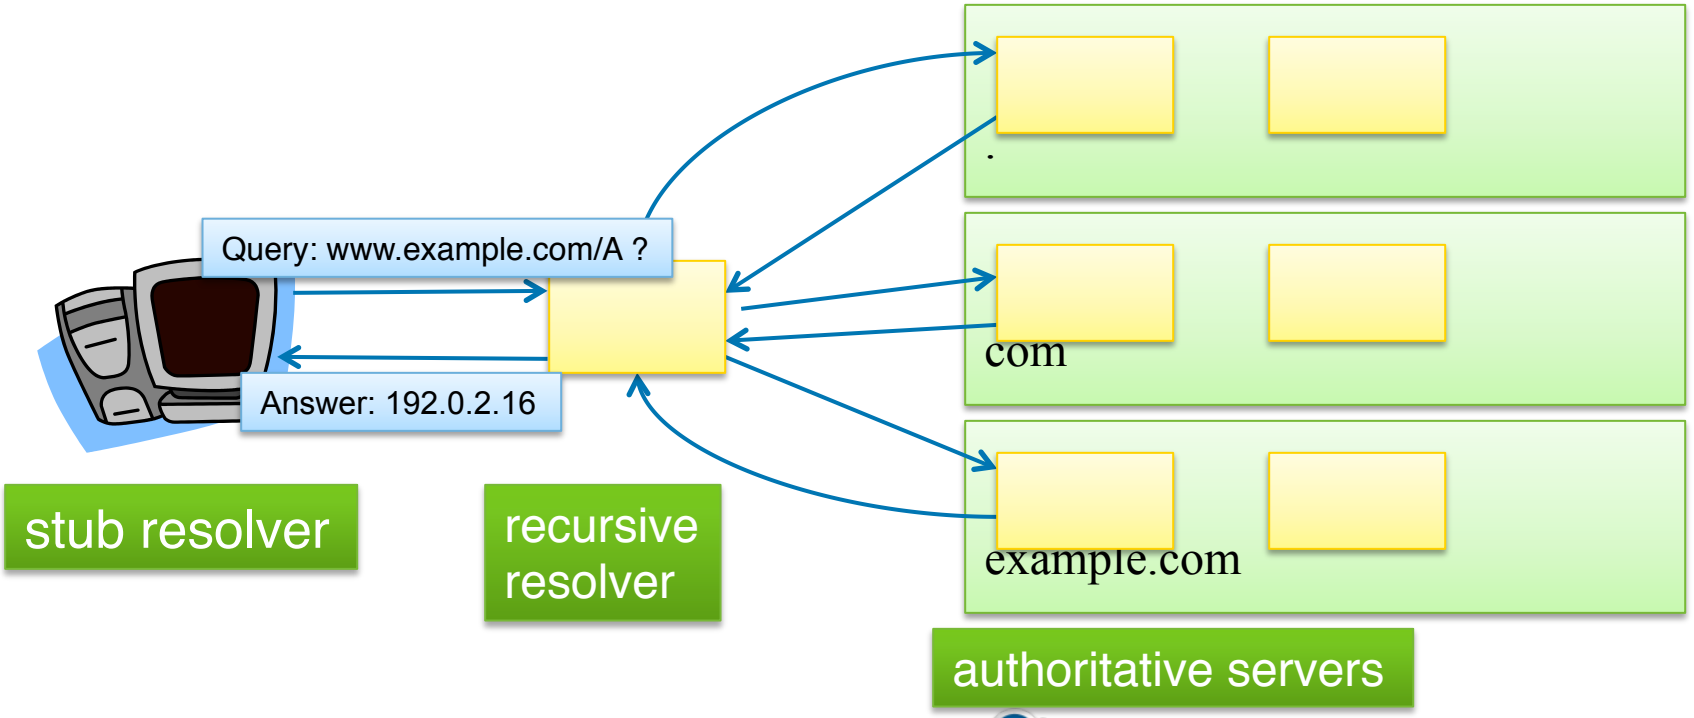

DNS Server Responsiveness

- At least one server must be responsive for a given zone
- Most resolver implementations prefer servers with lower response times (RFC 1035)
- Query response time for stub resolver is largely based on:
  - 1) Contents of cache of recursive resolver(s)
  - 2) Response time from authoritative server(s)

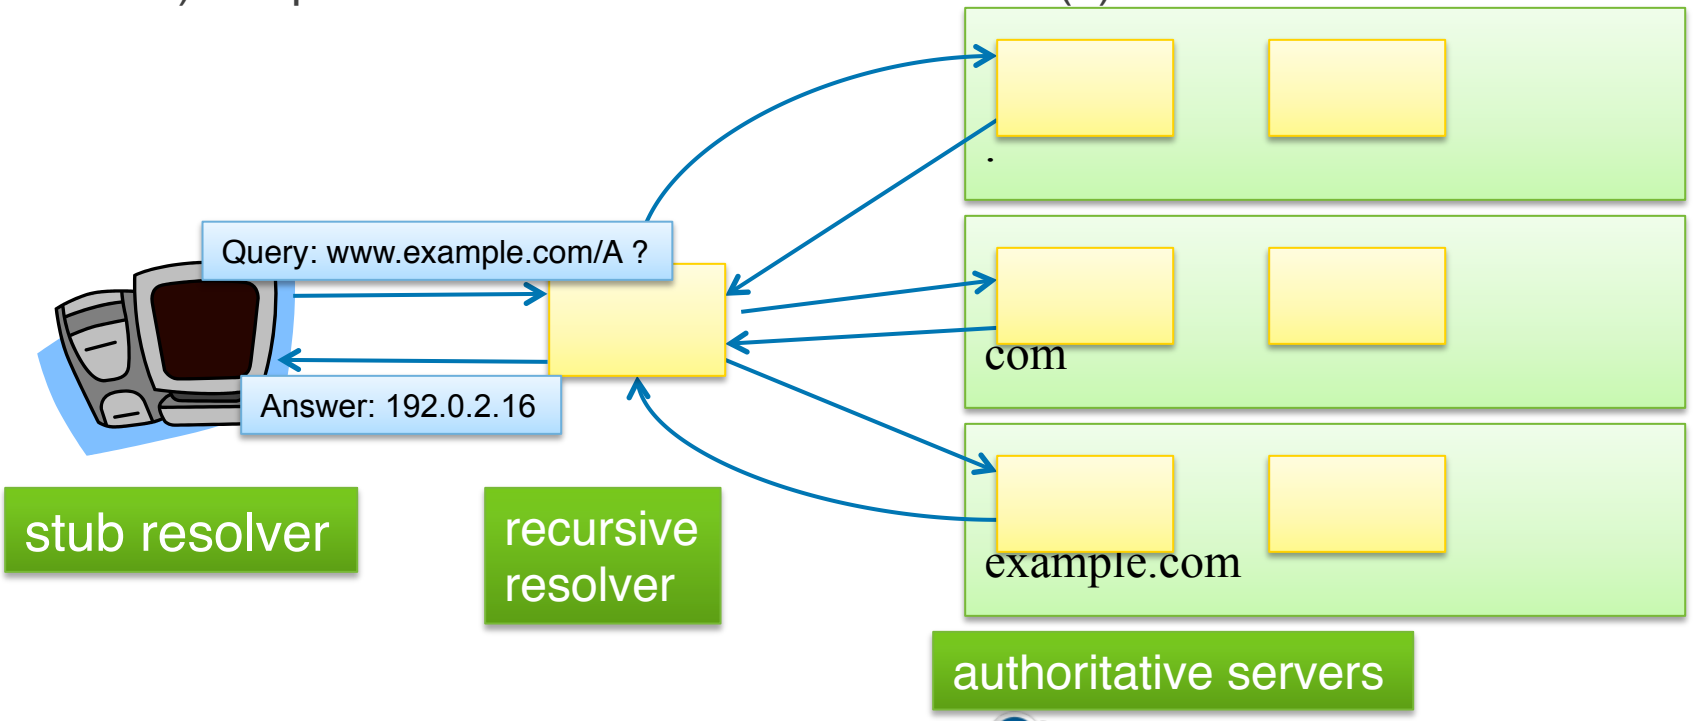

Anycast

- Many root and TLD servers employ anycast
- Different server instances respond from different autonomous systems for same address
- Queries from clients (recursive resolvers) are routed to closest anycast instance

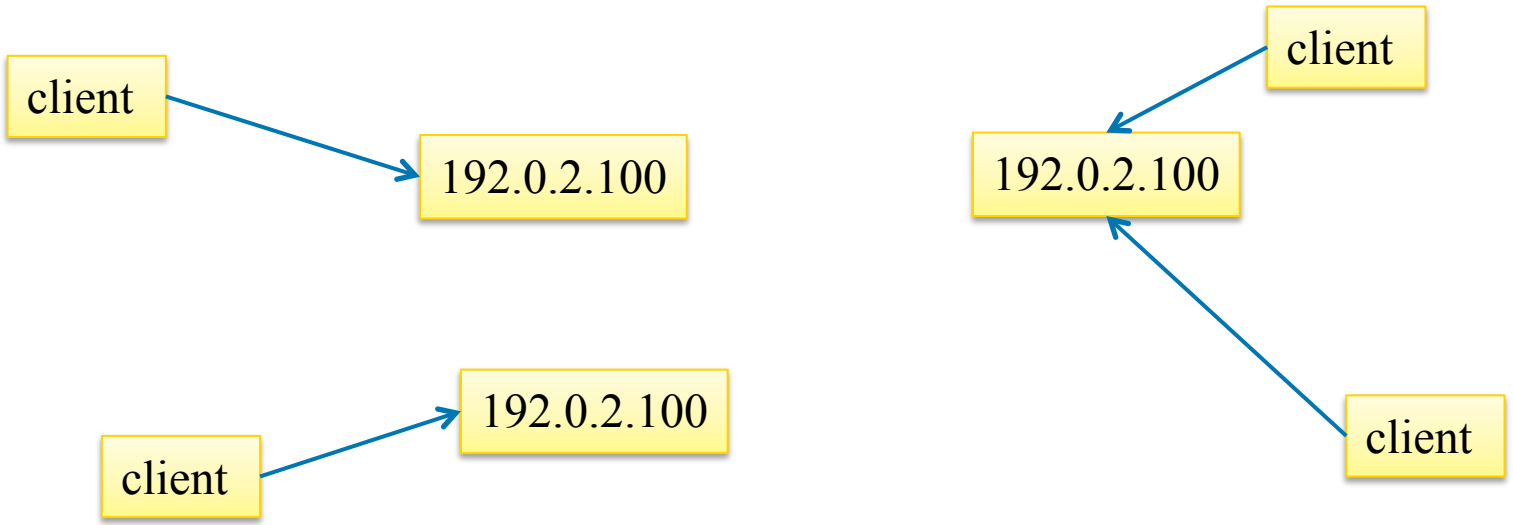

Multiple Paths

- clients x servers x IP versions x anycast instances =
  - Diverse paths
  - Many middleboxes
  - Non-determinism
  - Potential for variable client experience

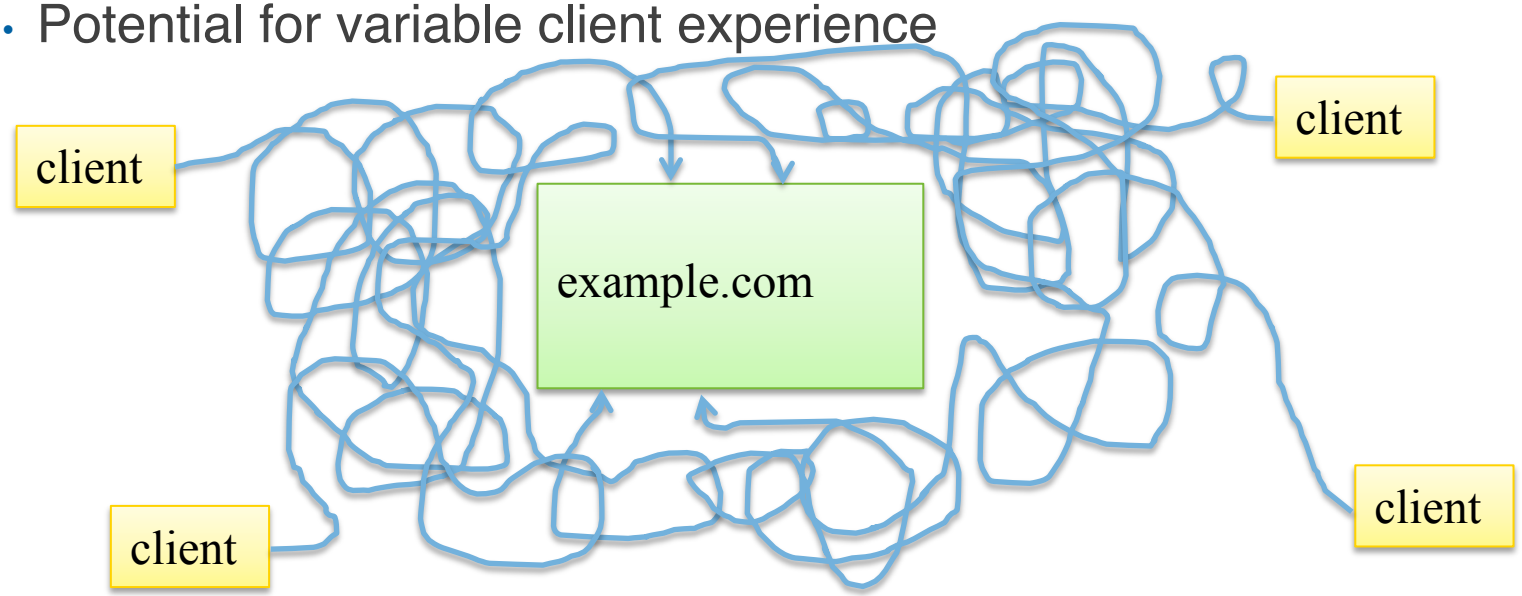

EDNS, DNSSEC, and Response Sizes

- EDNS (extended DNS) enables larger (> 512-byte) DNS UDP responses
- DNSSEC adds special records to the DNS, including public keys and cryptographic signatures
  - Requires EDNS
  - Results in a general increase in DNS response size
- Some middle boxes mishandle EDNS/DNSSEC
  - Drop EDNS requests/responses
  - Strip EDNS/DNSSEC records from requests/responses
  - Drop/mishandle IP fragments

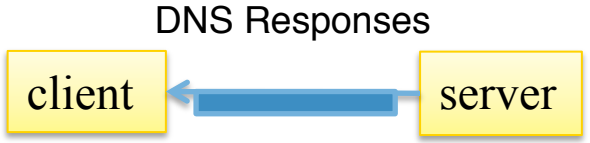

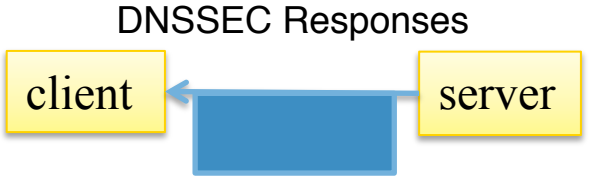

DNS Resolver Middlebox Workarounds

- Resolvers lower UDP max payload size (EDNS option) on timeout
  - Result: overcome path maximum transmission unit (PMTU) bottleneck
  - Side-effects:
    - PMTU problems masked by resolver workarounds
    - Additional RTTs and (sometimes) forced TCP usage
- Resolvers avoid sending EDNS packets to/through non-EDNS-compatible servers/paths
  - Result: Get an answer from otherwise unresponsive servers
  - Side-effect: DNSSEC records not retrievable from affected servers

Multi-perspective DNS Measurement

DNS Analysis Using DNSViz (dnsget command line)

- Online analysis (query/response) of DNS name and servers
- Output: Serialized (JSON) DNS analysis, including query/response diagnostics (timeout retries, reduced payload, EDNS disabling)

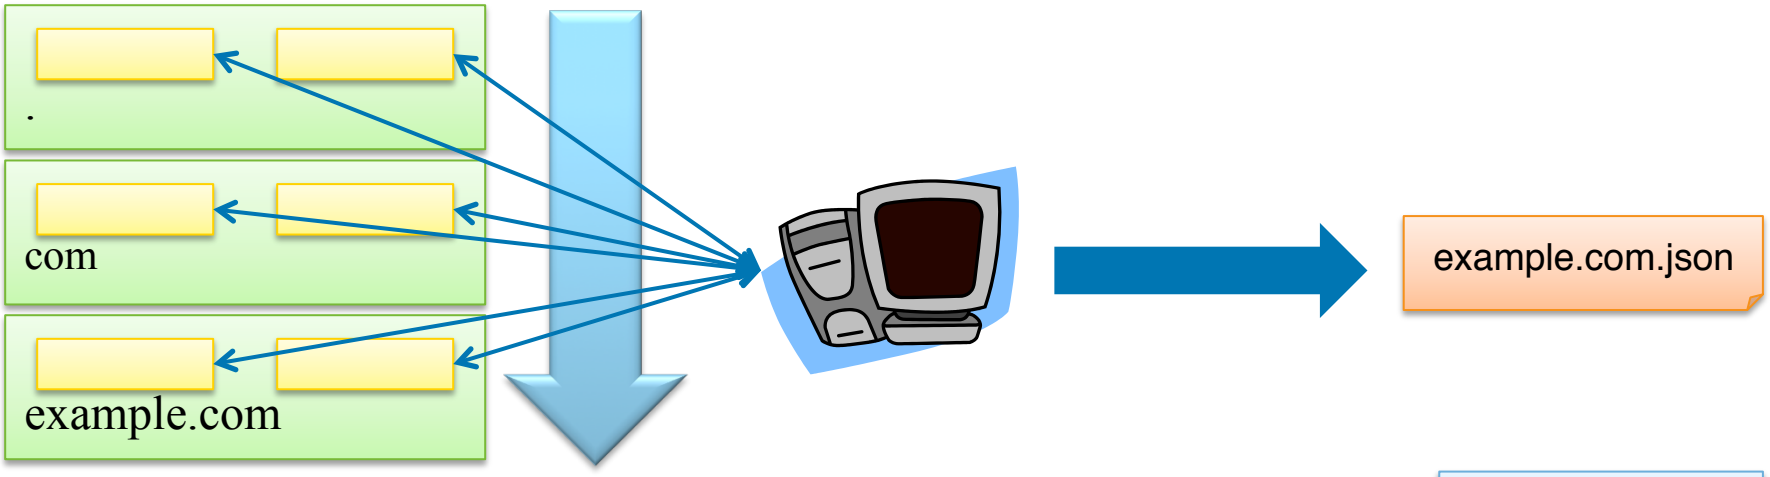

Distributed root/TLD Measurement Using CAIDA Ark Nodes

- DNSViz code installed on Ark nodes
  - 32 nodes (FreeBSD)
  - 27 countries
- Queries: NS/SOA/DNSKEY/DS, NXDOMAIN/NODATA
- Transport/Network: TCP, UDP, IPv4, IPv6
- Time: 4x daily for six days

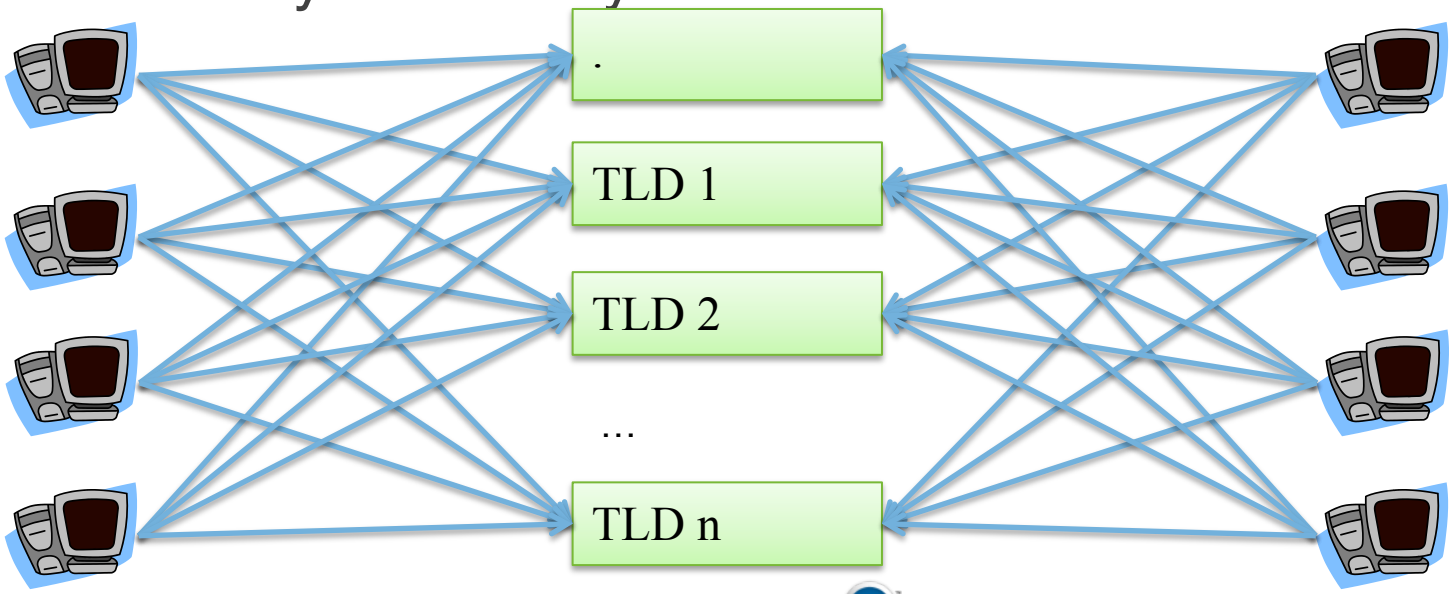

Server Responsiveness – IPv4

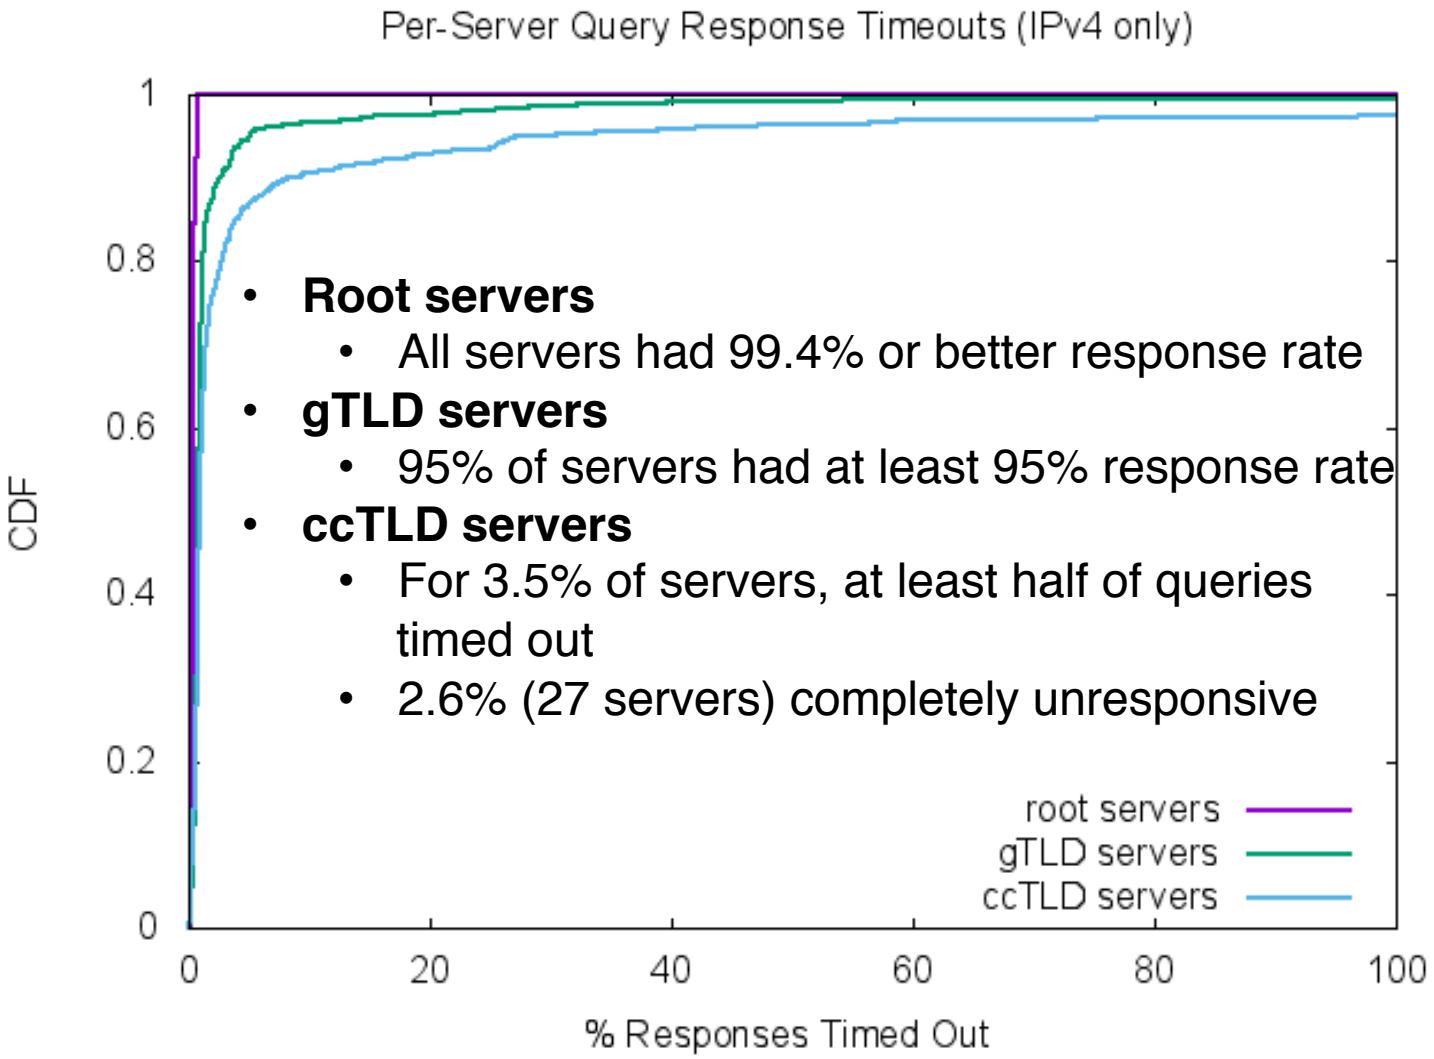

- **Root servers**
  - All servers had 99.4% or better response rate
- **gTLD servers**
  - 95% of servers had at least 95% response rate
- **ccTLD servers**
  - For 3.5% of servers, at least half of queries timed out
  - 2.6% (27 servers) completely unresponsive

Root Server Responsiveness per Client – IPv4

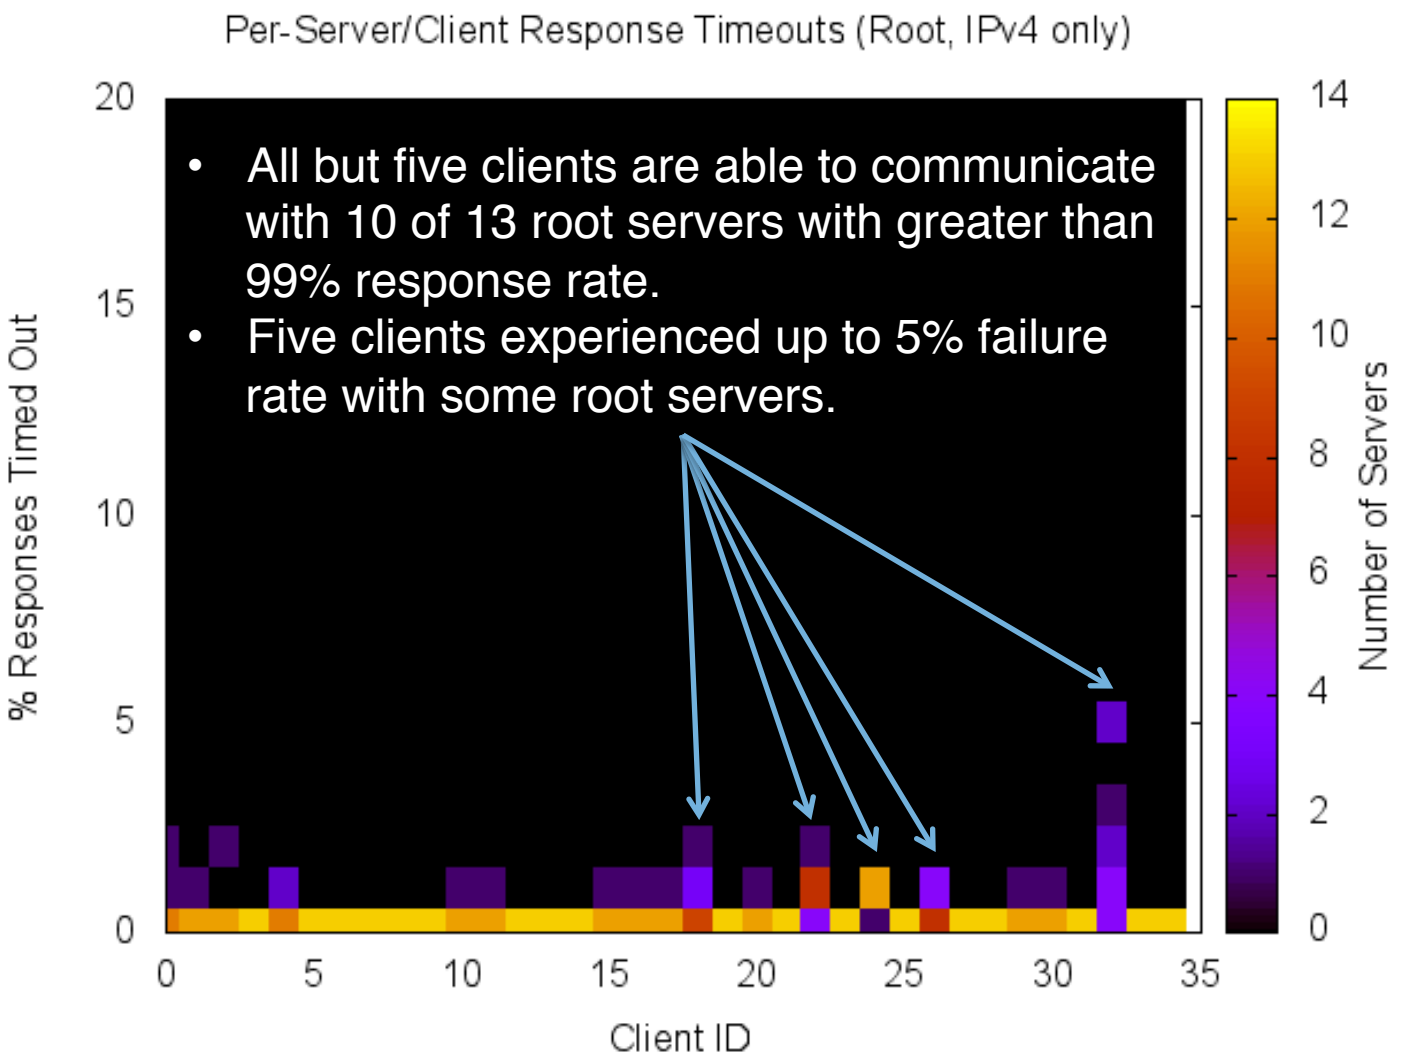

gTLD Server Responsiveness per Client – IPv4

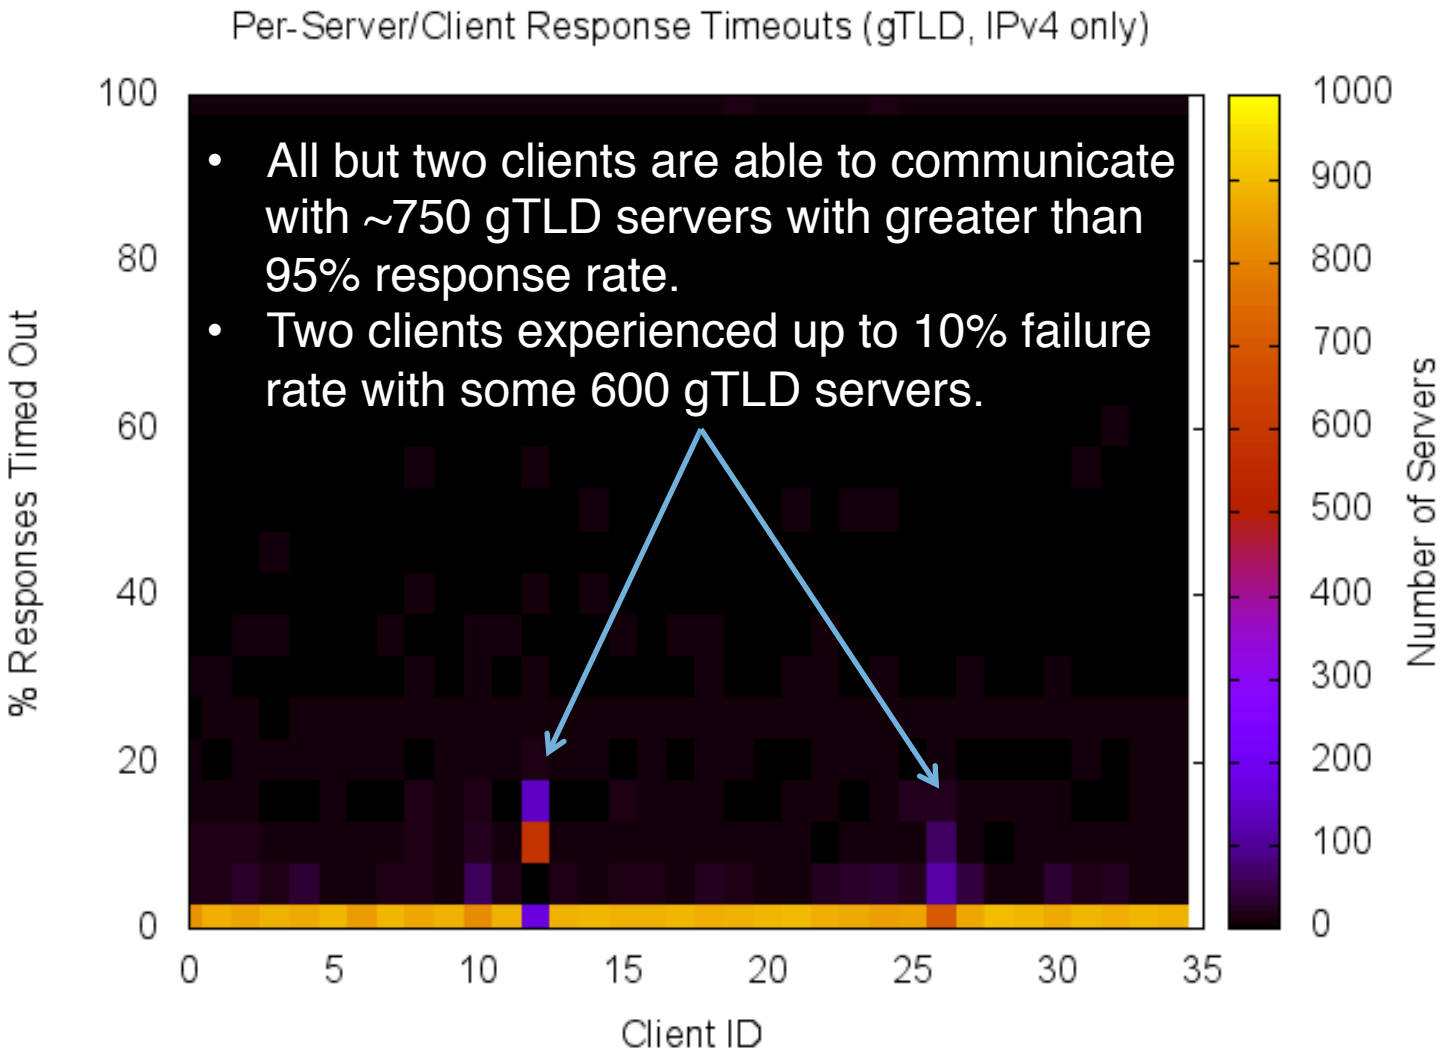

ccTLD Server Responsiveness per Client – IPv4

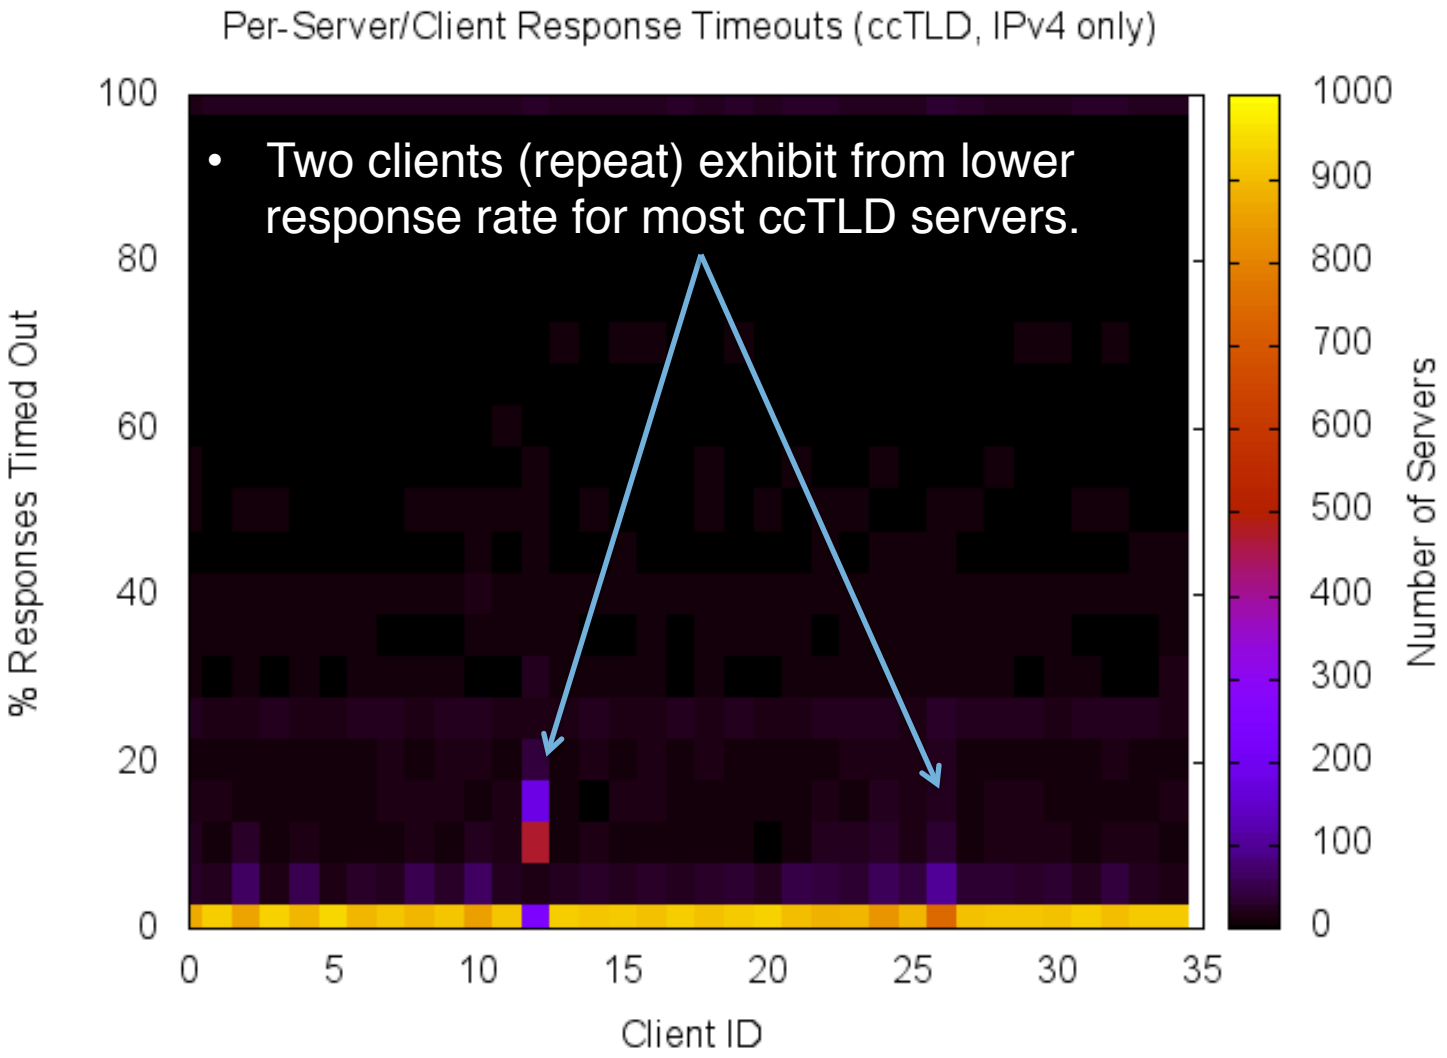

Server Responsiveness – IPv6

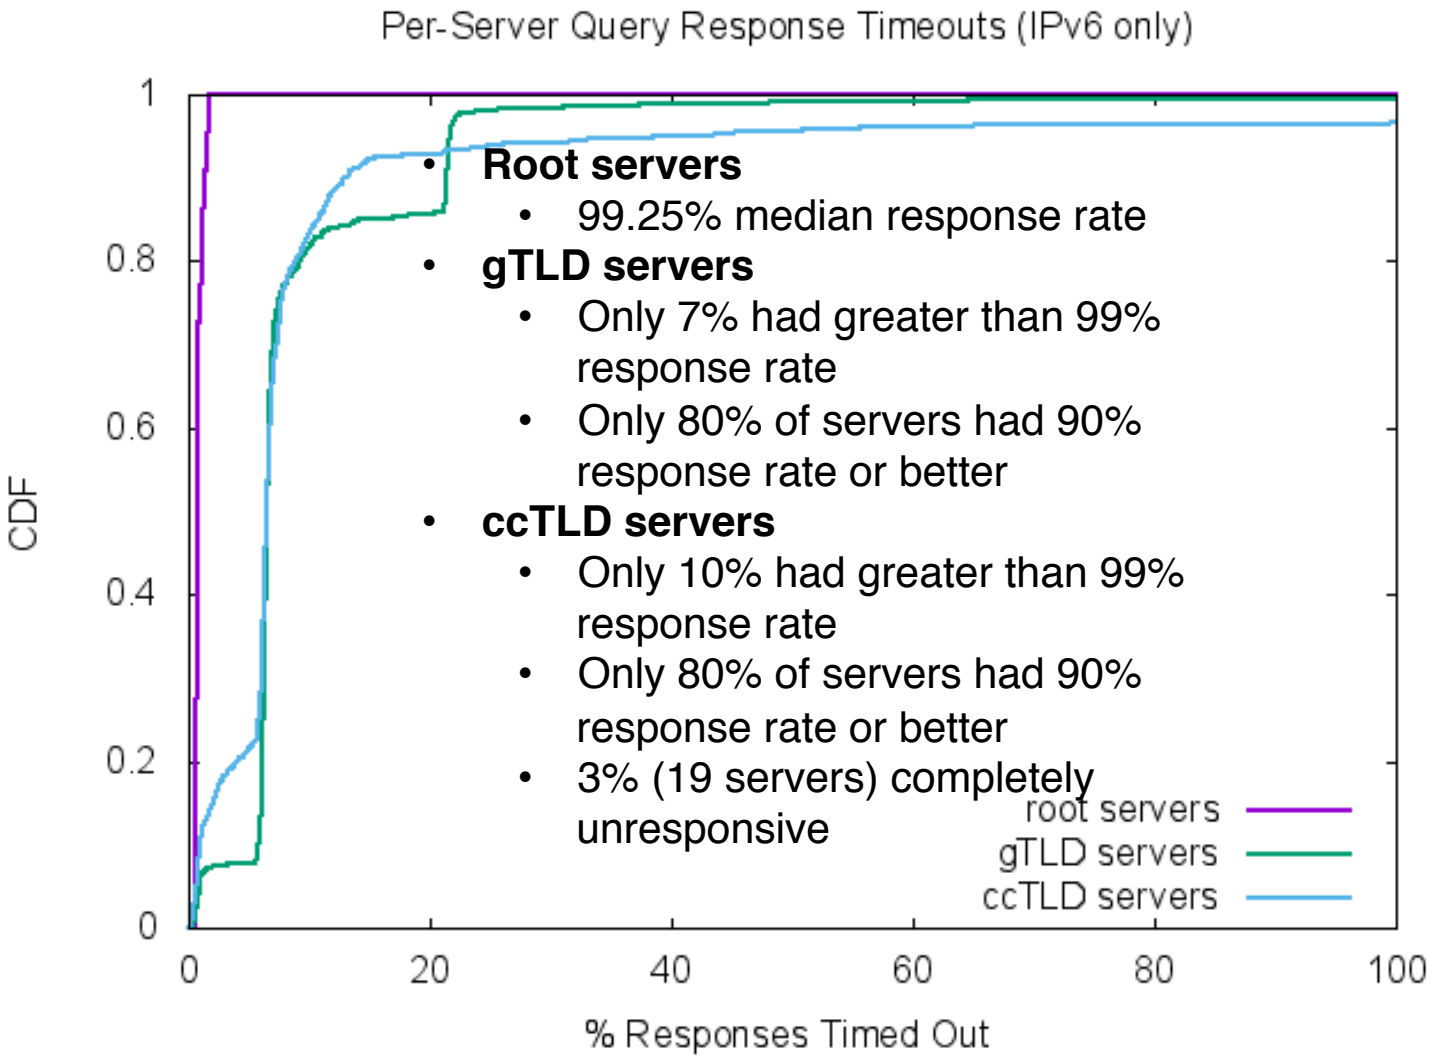

- **Root servers**
  - 99.25% median response rate
- **gTLD servers**
  - Only 7% had greater than 99% response rate
  - Only 80% of servers had 90% response rate or better
- **ccTLD servers**
  - Only 10% had greater than 99% response rate
  - Only 80% of servers had 90% response rate or better
  - 3% (19 servers) completely unresponsive

Root Server Responsiveness per Client – IPv6

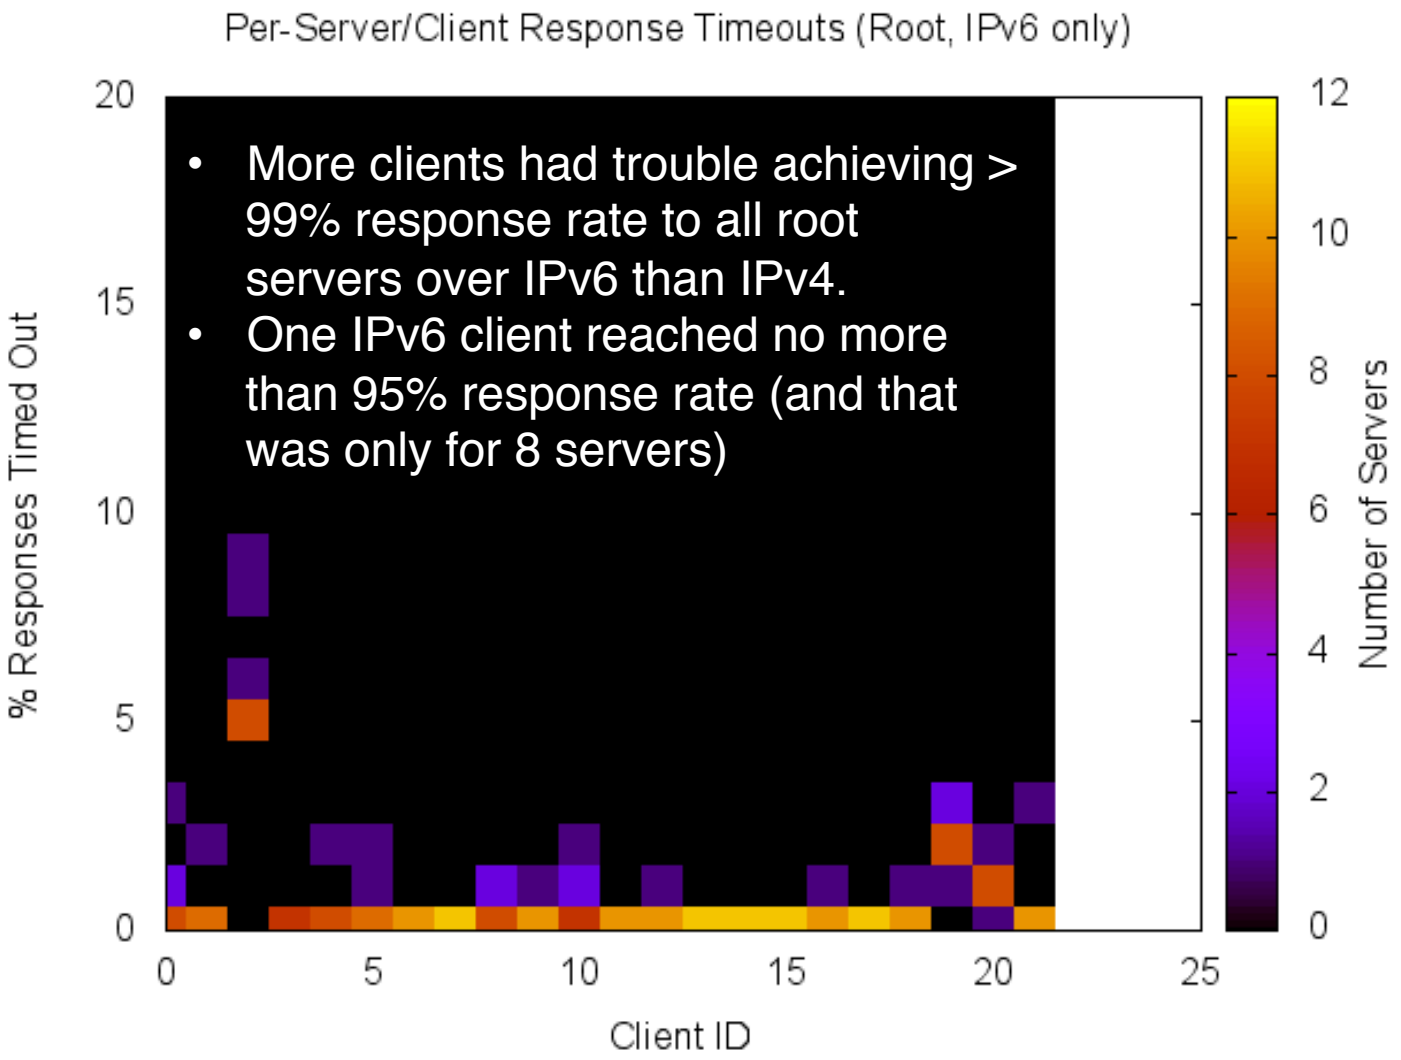

gTLD Server Responsiveness per Client – IPv6

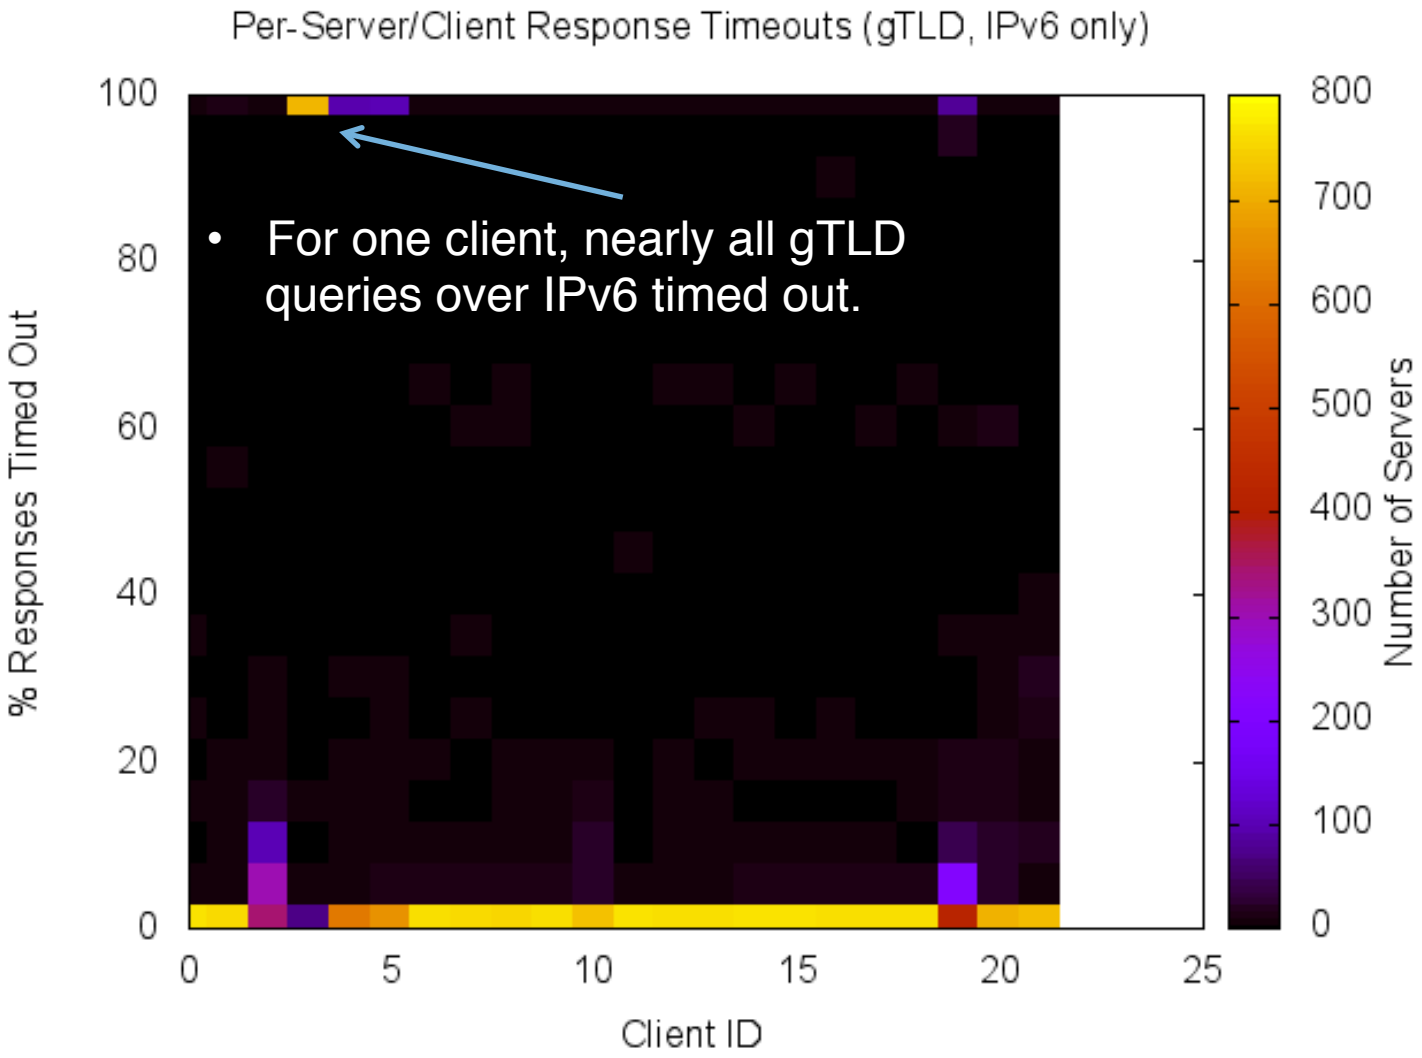

ccTLD Server Responsiveness per Client – IPv6

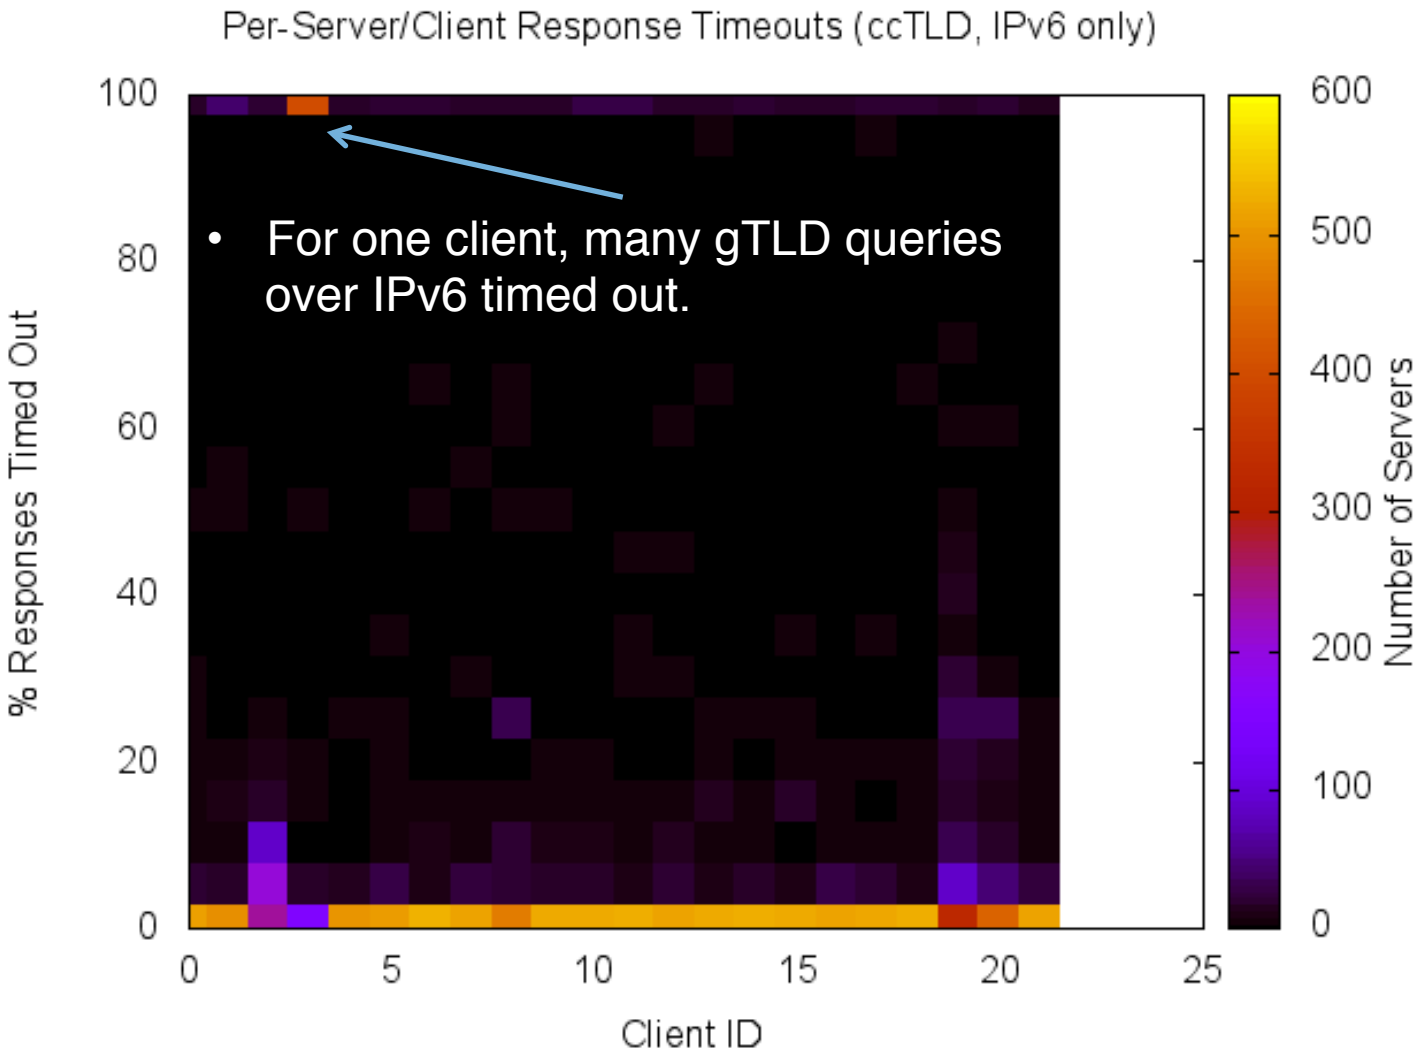

Server Responsiveness – IPv6 (without “client 2”)

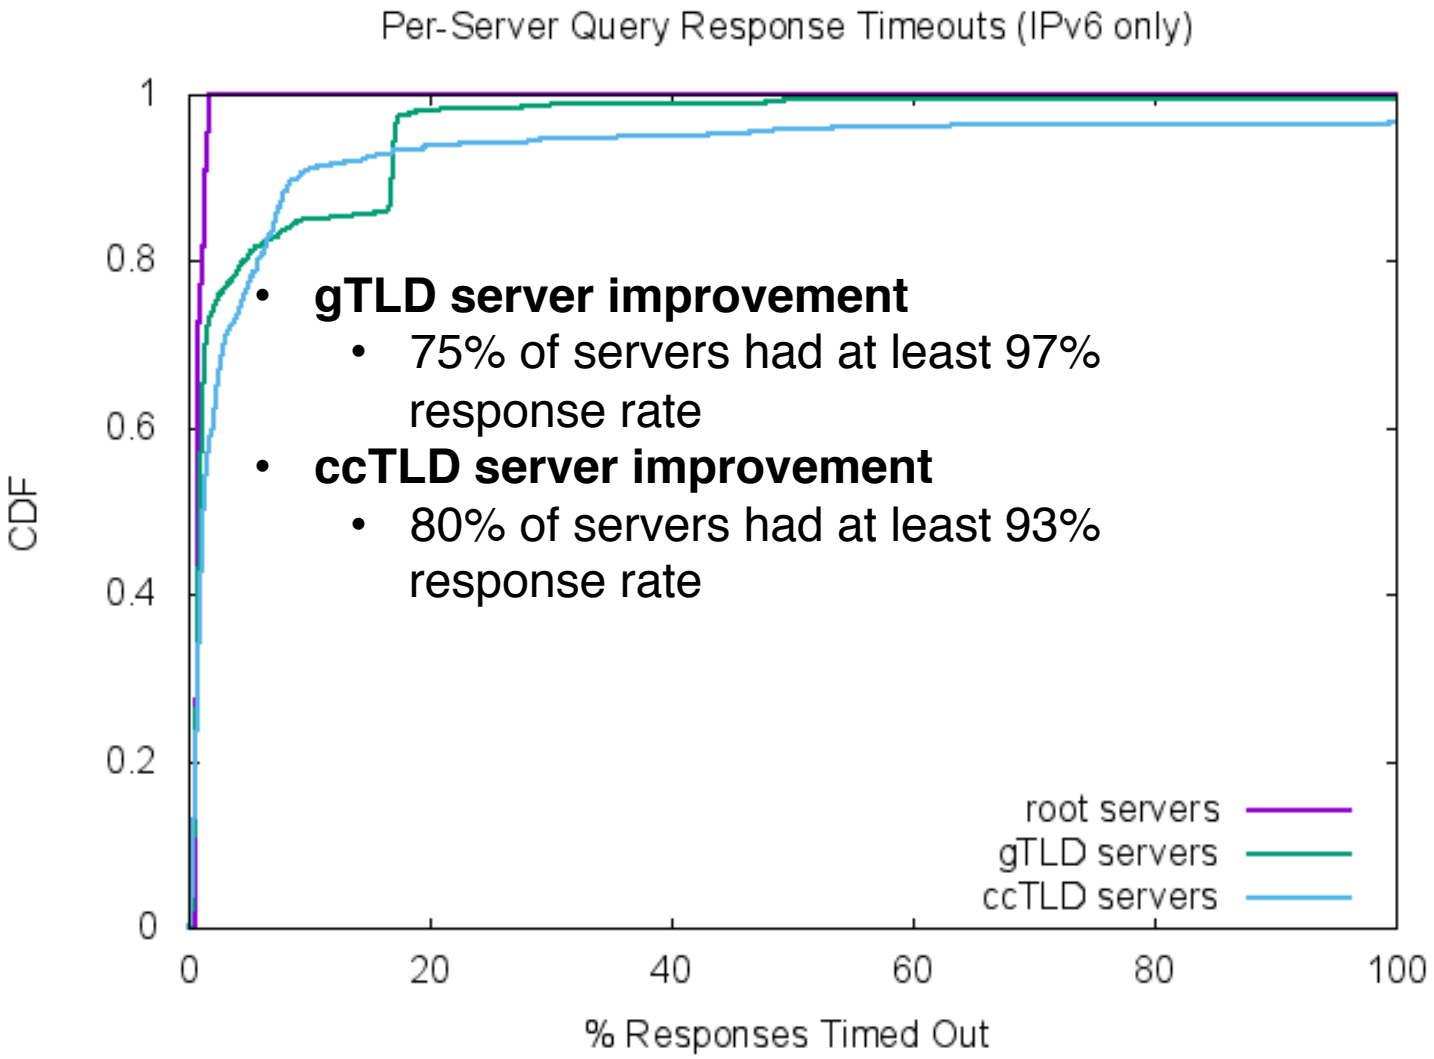

- **gTLD server improvement**
  - 75% of servers had at least 97% response rate
- **ccTLD server improvement**
  - 80% of servers had at least 93% response rate

Response Time – IPv4

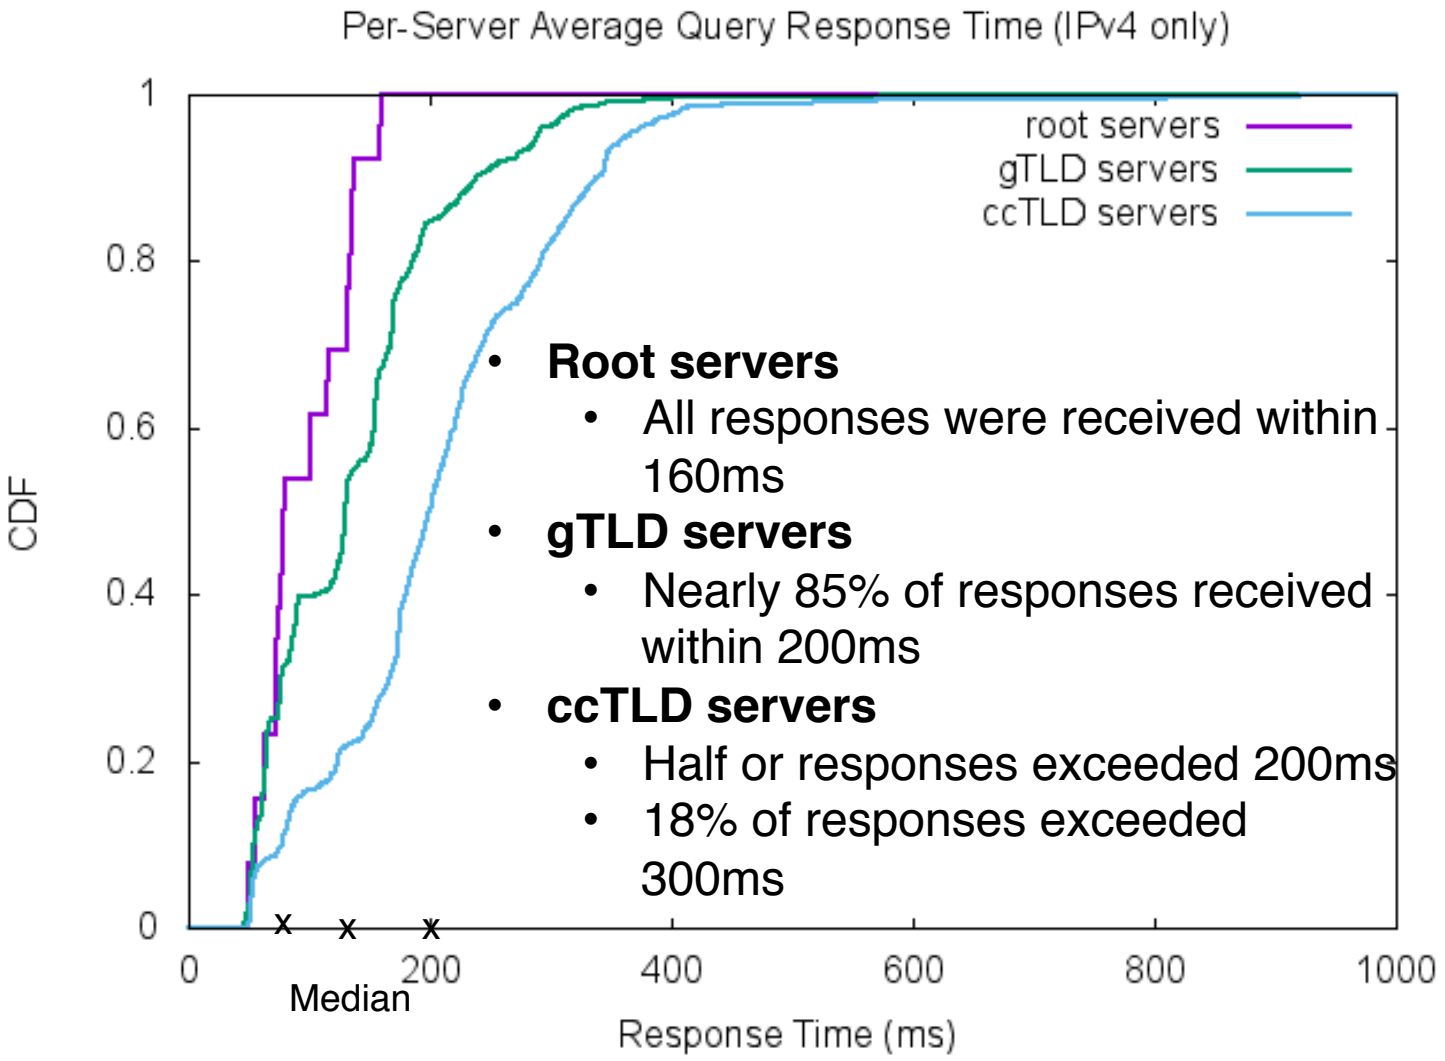

- **Root servers**
  - All responses were received within 160ms
- **gTLD servers**
  - Nearly 85% of responses received within 200ms
- **ccTLD servers**
  - Half or responses exceeded 200ms
  - 18% of responses exceeded 300ms

Response Time – IPv6

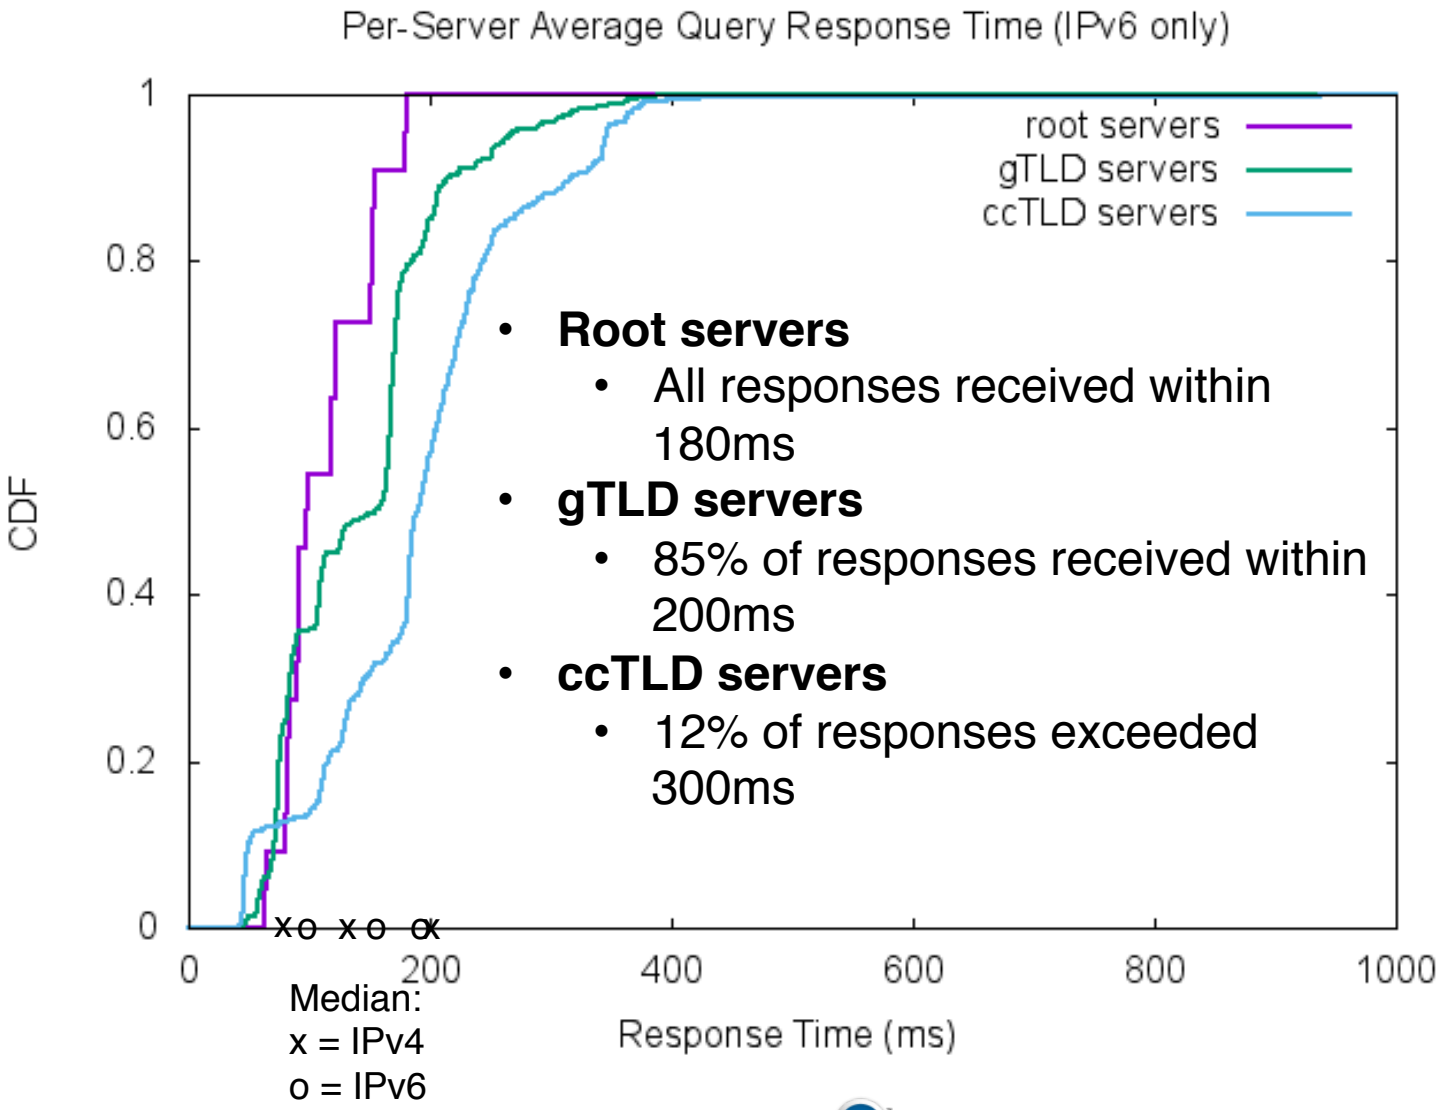

- **Root servers**
  - All responses received within 180ms
- **gTLD servers**
  - 85% of responses received within 200ms
- **ccTLD servers**
  - 12% of responses exceeded 300ms

Root Server Response Time per Client – IPv4

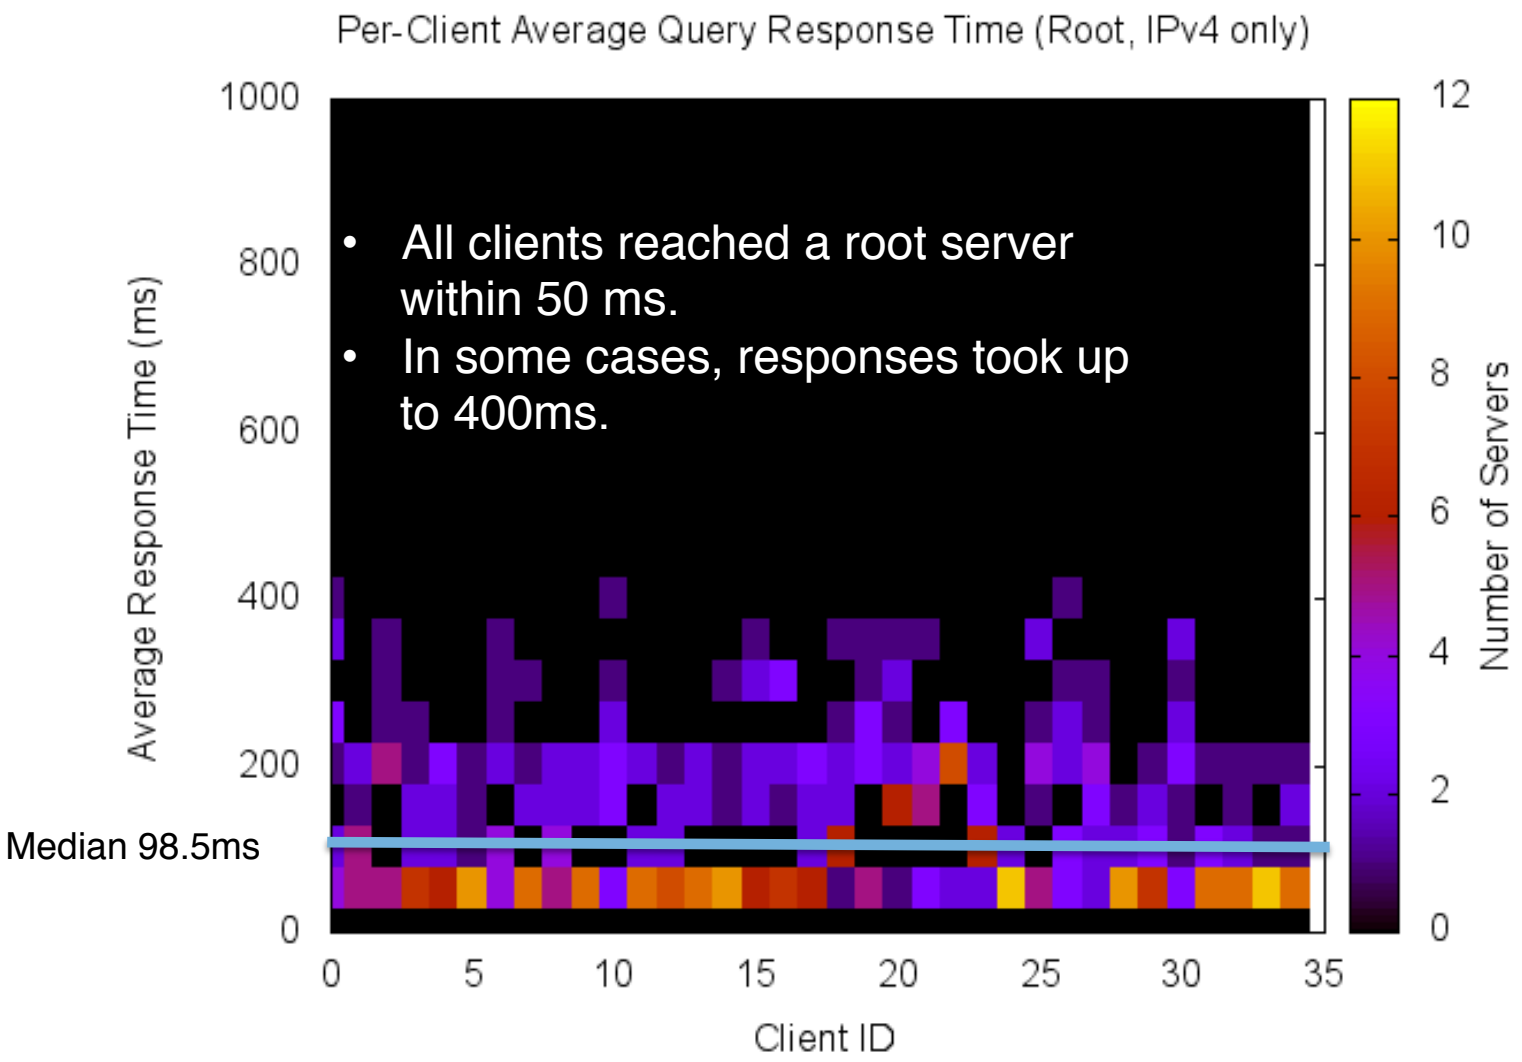

Root Server Response Time per Client – IPv6

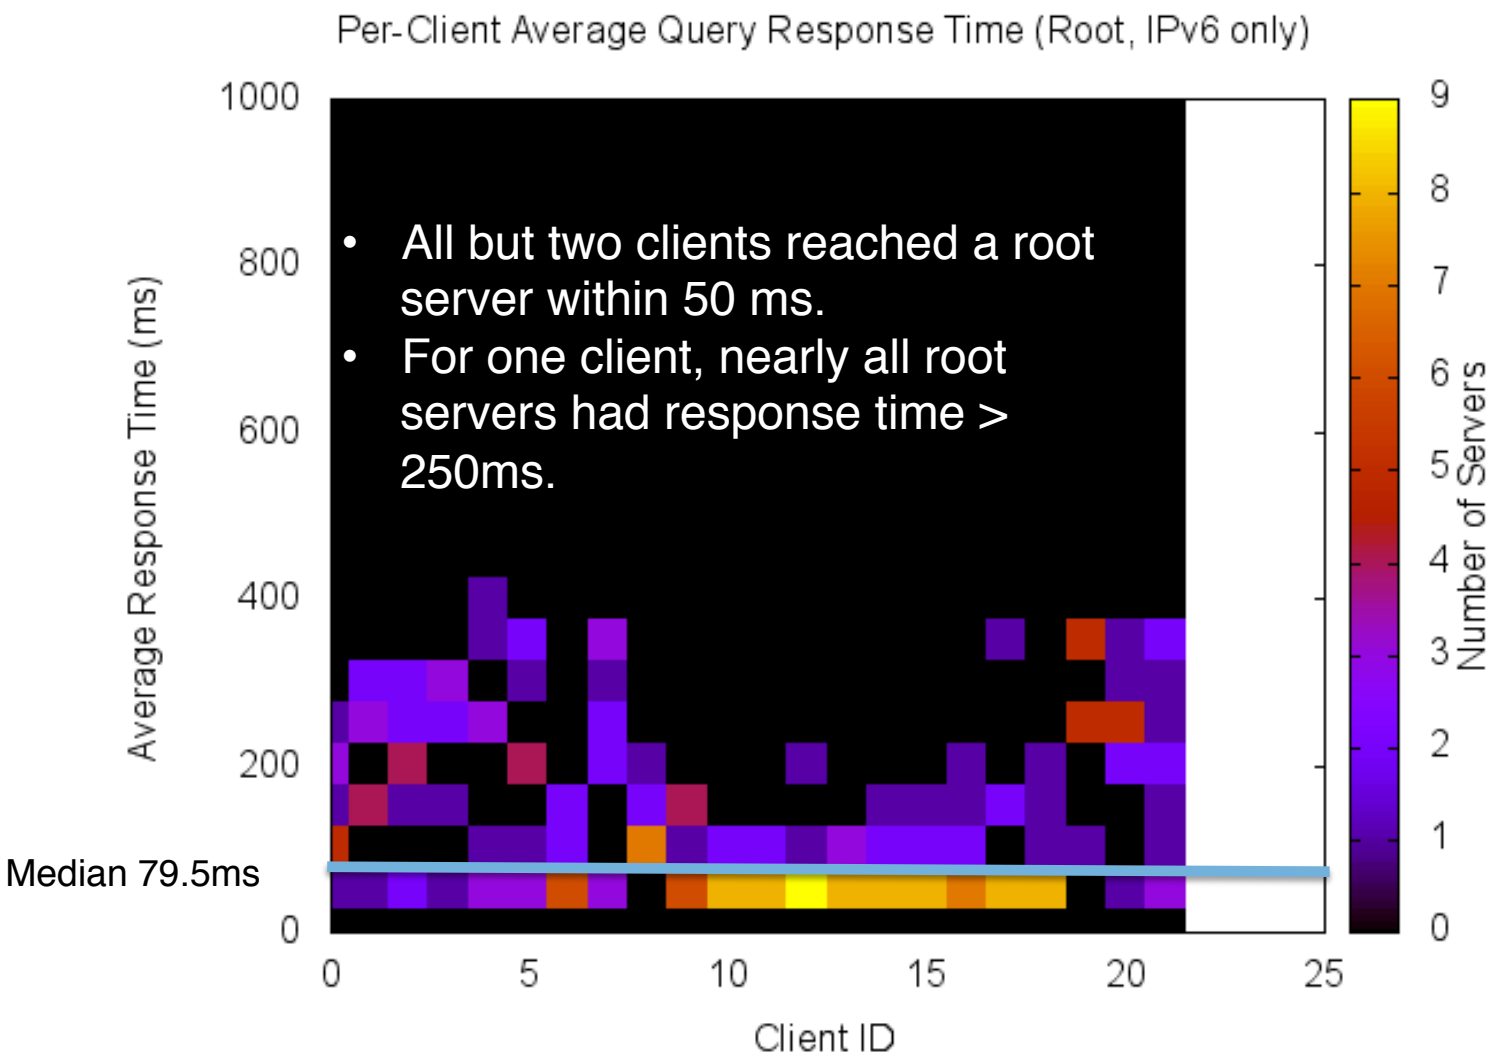

gTLD Server Response Time per Client – IPv4

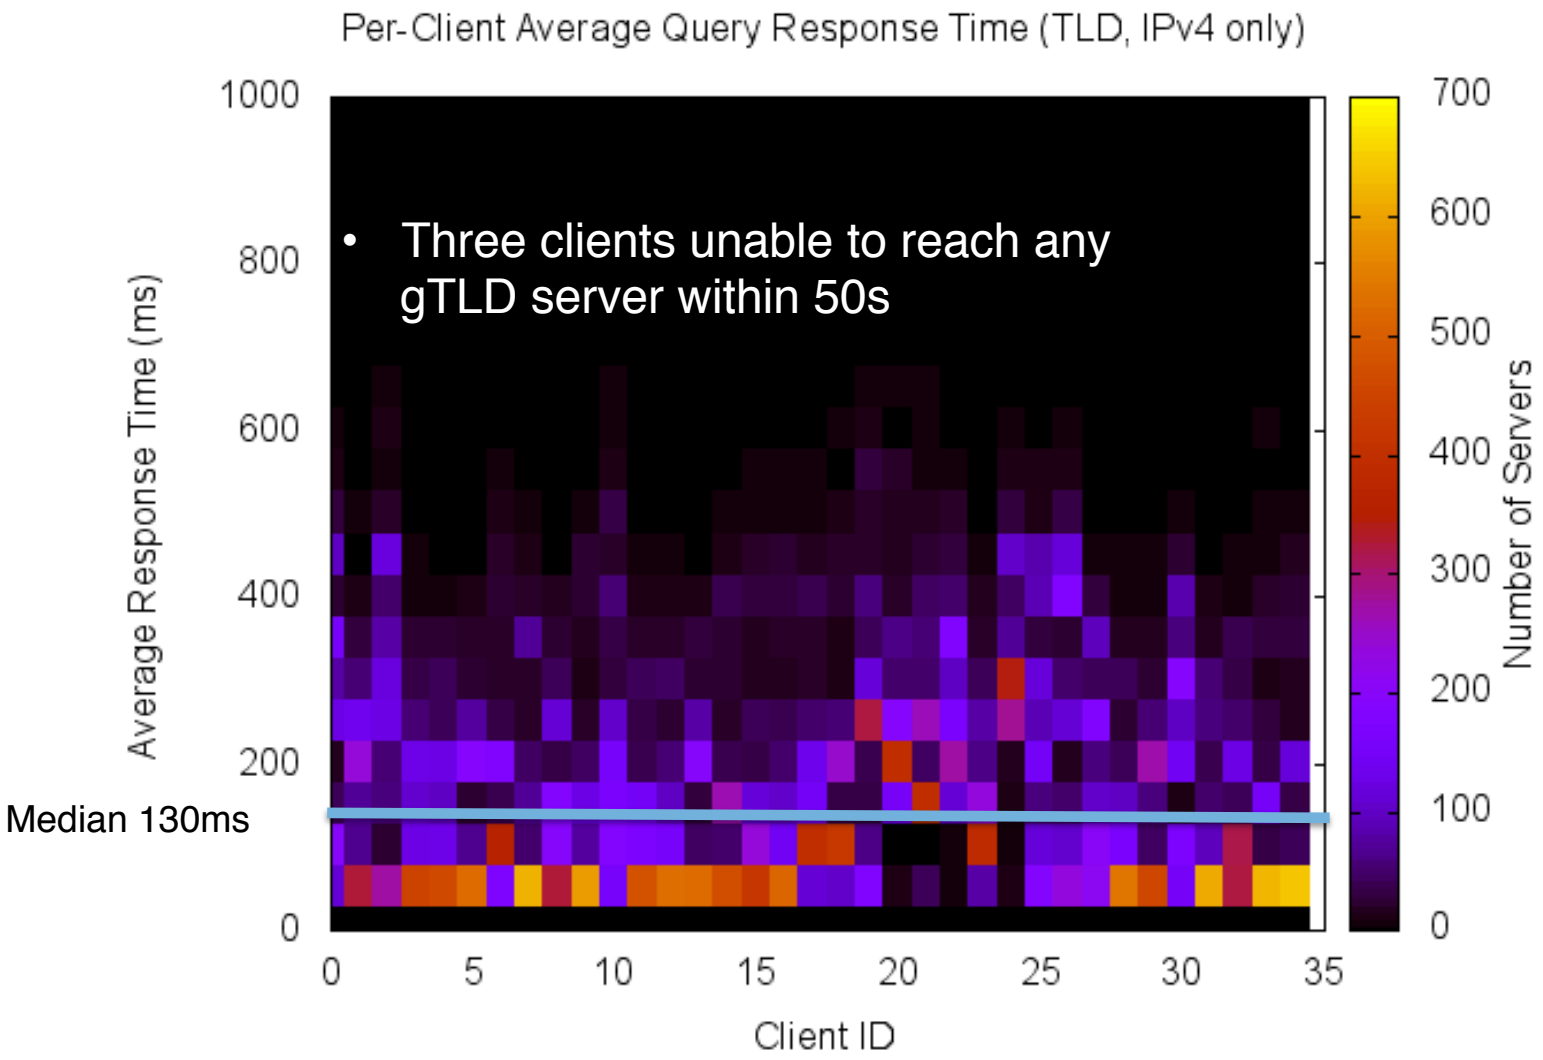

gTLD Server Response Time per Client – IPv6

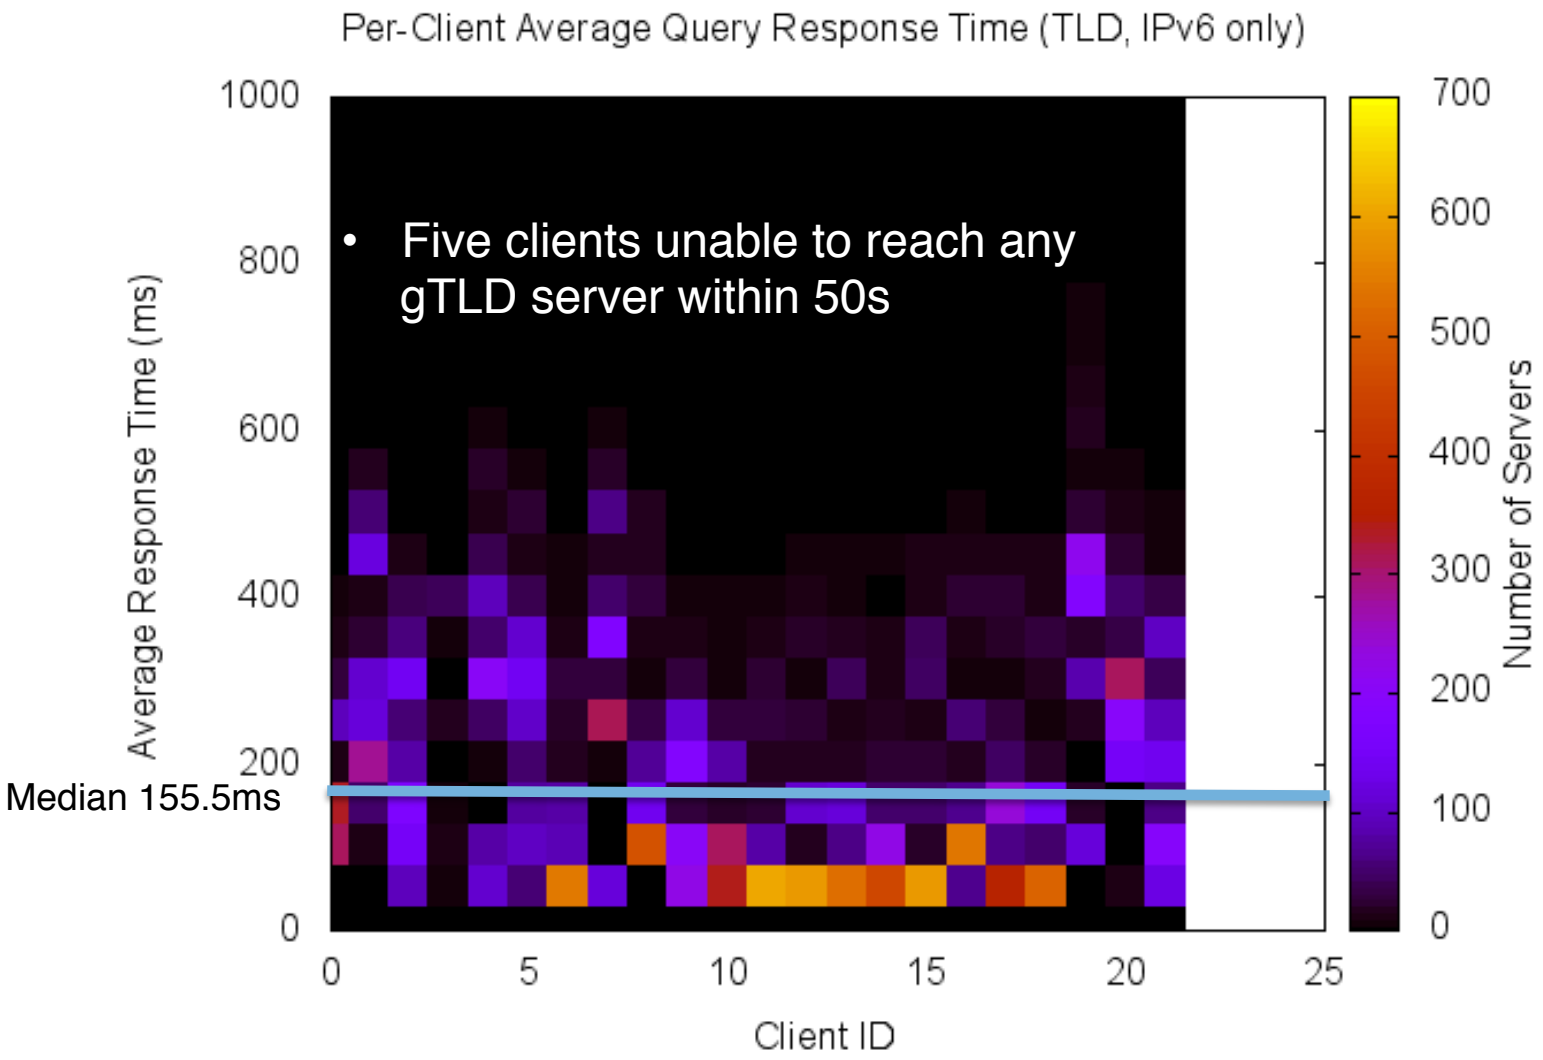

ccTLD Server Response Time per Client – IPv4

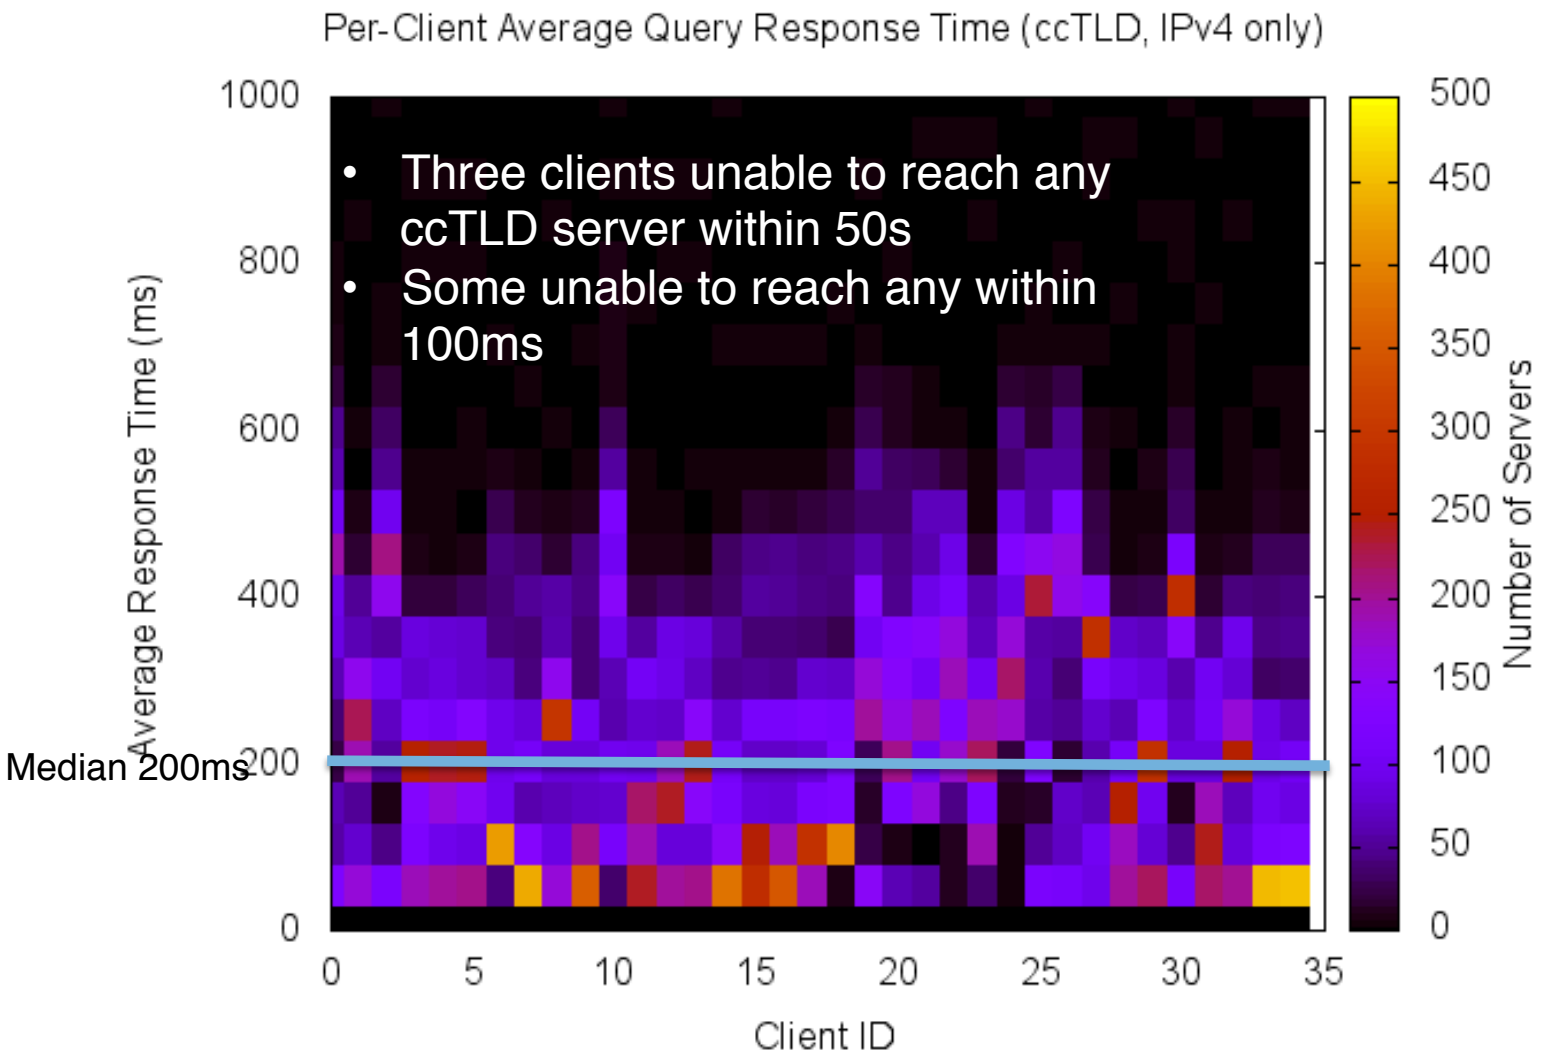

ccTLD Server Response Time per Client – IPv6

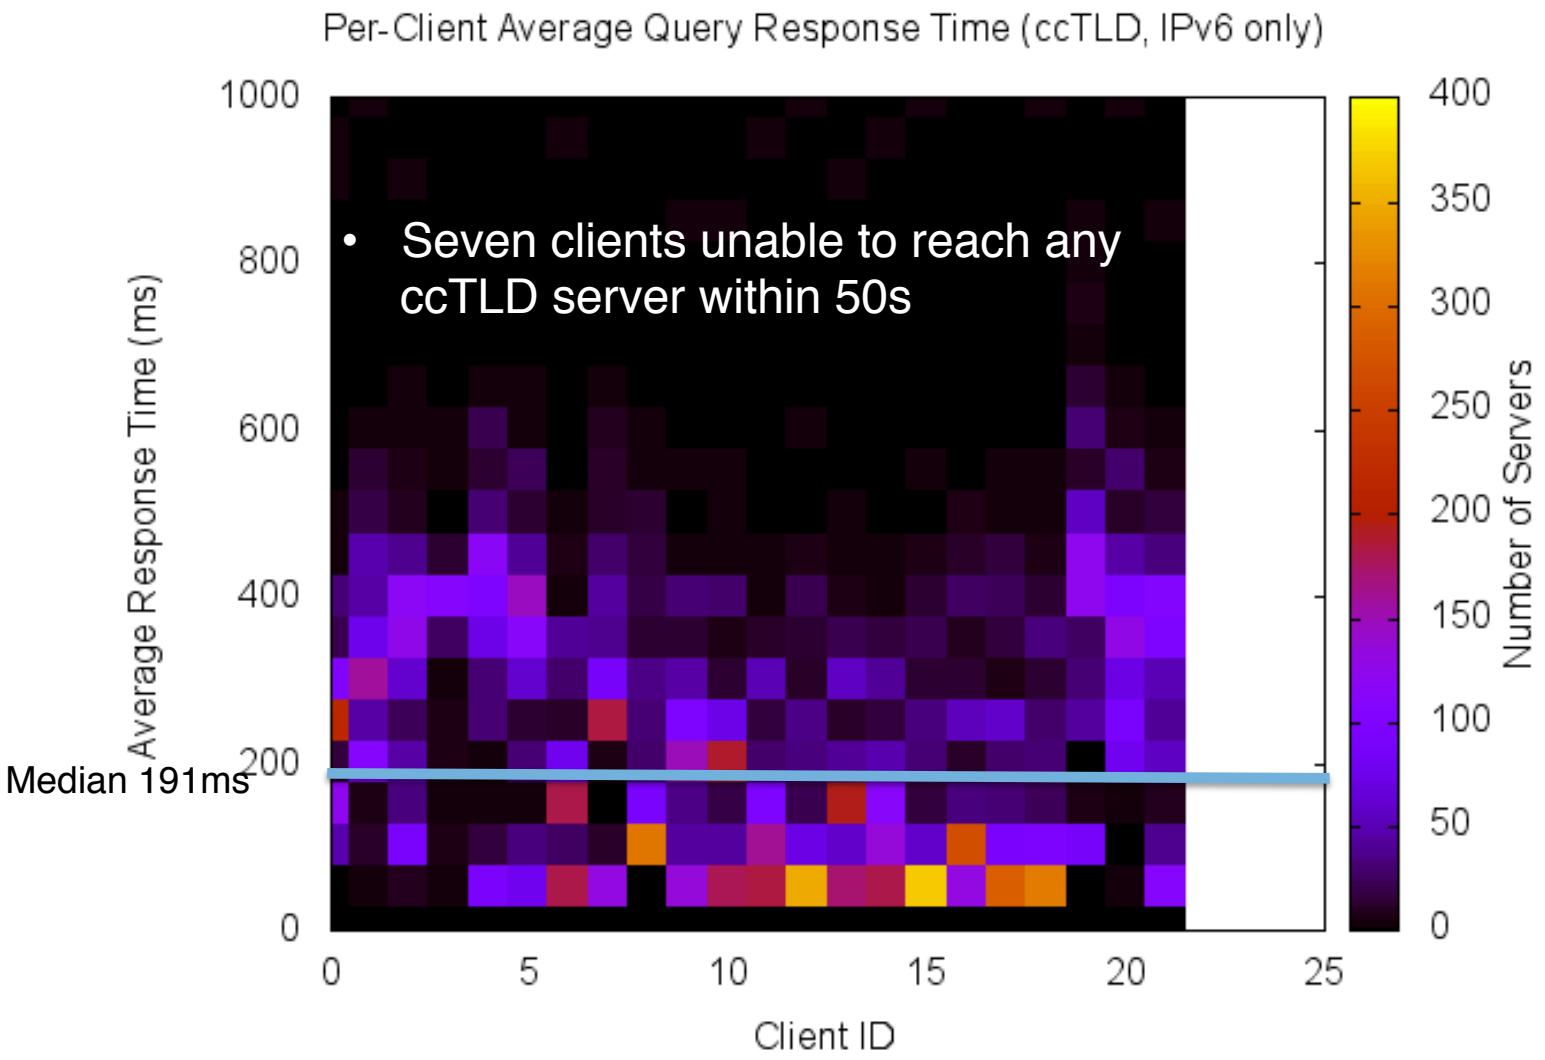

Summary

- DNS name resolution paths can be diverse.
- A multi-perspective analysis can help understand general resolver experience.
- Results from preliminary experimentation:
  - Root server communication is generally quick and stable from all instrumented locations.
  - Most ccTLD/gTLD servers have reasonable response rates and response times.
  - Some (ccTLD) servers are not available from any vantage point.
  - Response times from root are generally lower than those from gTLD/ccTLD servers.
  - Median IPv6 response time from ccTLD servers is less than median IPv4 response time.

Future Work

- Further analyze/refine preliminary data/methodologies
- Analyze path similarity between clients/servers
- Identify EDNS/PMTU issues between clients/servers
- Quantify impact of response rate/response time

powered by

VERISIGN™