# Content Distribution Networks

COS 518: Advanced Computer Systems
Lecture 18

Mike Freedman

## Content Distribution Network

- Proactive content replication
  - Content provider (e.g., CNN) contracts with a CDN
- CDN replicates the content
  - On many servers spread throughout the Internet
- Updating the replicas
  - Updates pushed to replicas when the content changes

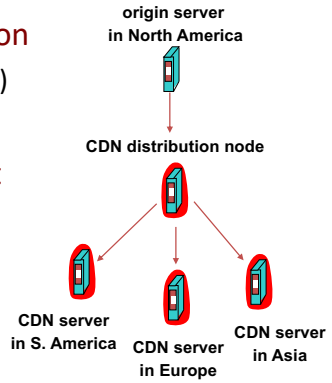




## Server Selection Policy

- Live server
  - For availability
- Lowest load
  - To balance load across the servers
- Closest
  - Nearest geographically, or in round-trip time
- Best performance
  - Throughput, latency, …
- Cheapest bandwidth, electricity, …

**Requires continuous monitoring of liveness, load, and performance**



## Server Selection Mechanism

- Application
  - HTTP redirection

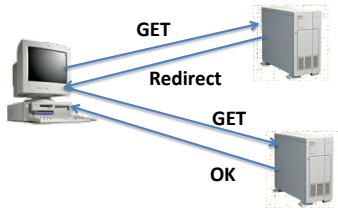


- Advantages
  - Fine-grain control
  - Selection based on client IP address
- Disadvantages
  - Extra round-trips for TCP connection to server
  - Overhead on the server

## Server Selection Mechanism

- Routing
  - Anycast routing
- Advantages
  - No extra round trips
  - Route to nearby server
- Disadvantages
  - Does not consider network or server load
  - Different packets may go to different servers
  - Used only for simple request-response apps

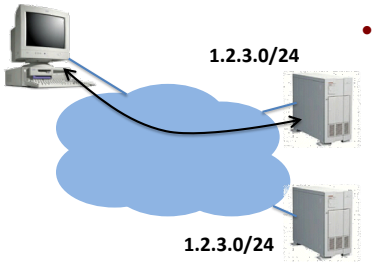


5

## Server Selection Mechanism

- Naming
  - DNS-based server selection

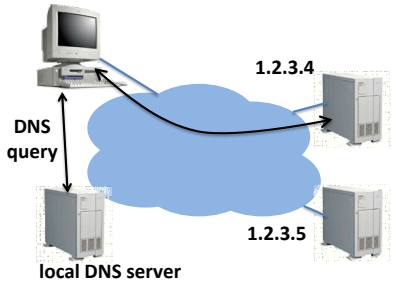


6

## A DNS lookup traverses DNS hierarchy

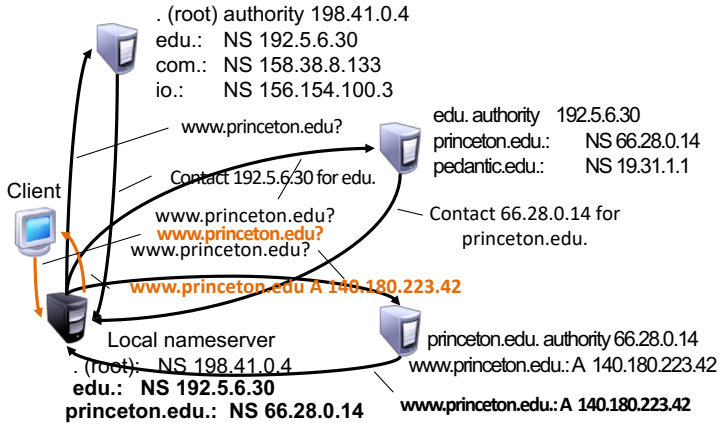


7

## DNS caching

- Performing all these queries takes time
  - And all this before actual communication takes place
- Caching can greatly reduce overhead
  - Top-level servers very rarely change, popular sites visited often
  - Local DNS server often has information cached
- How DNS caching works
  - All DNS servers **cache responses to queries**
  - Responses include a time-to-live (TTL) field, akin to cache expiry

8

## Server Selection Mechanism

- Naming
  - DNS-based server selection

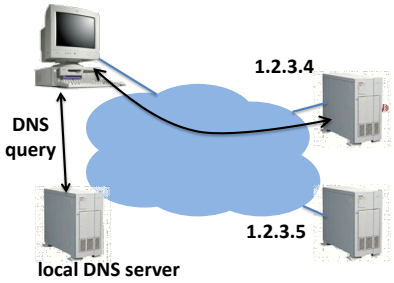


- Advantages
  - Avoid TCP set-up delay
  - DNS caching reduces overhead
  - Relatively fine control
- Disadvantage
  - Based on IP address of local DNS server
  - "Hidden load" effect
  - DNS TTL limits adaptation

9

## How Akamai Works

10

11

## How Akamai Uses DNS

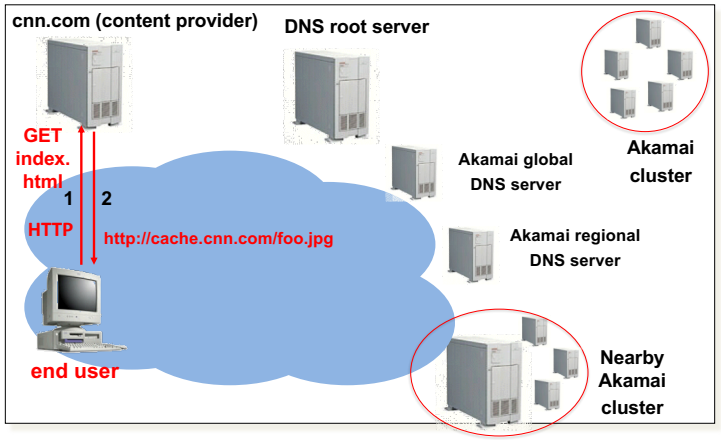


12

## How Akamai Uses DNS

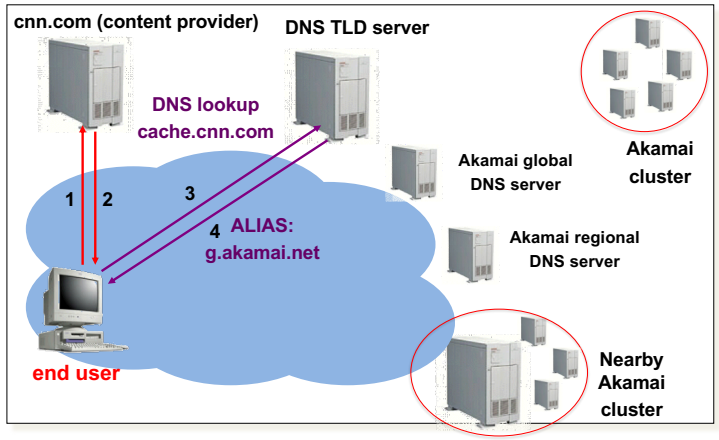

13

# How Akamai Uses DNS

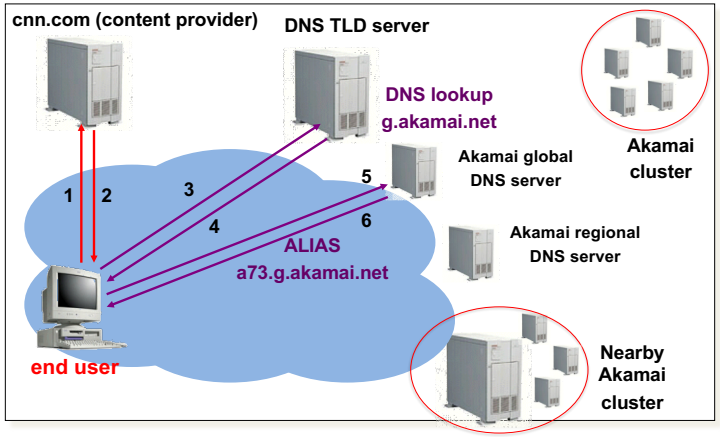


14

# How Akamai Uses DNS

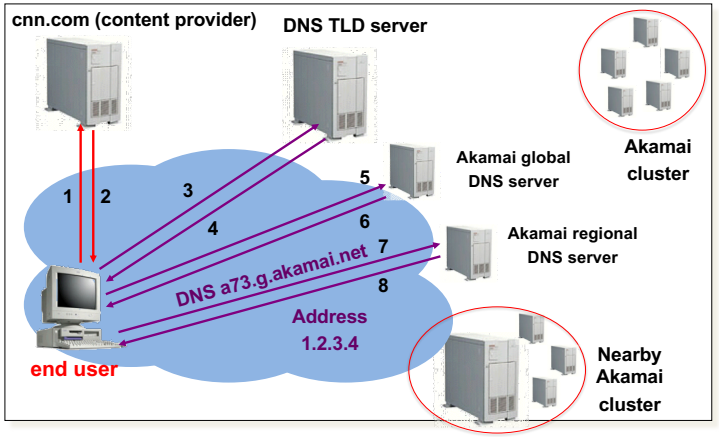


15

# How Akamai Uses DNS

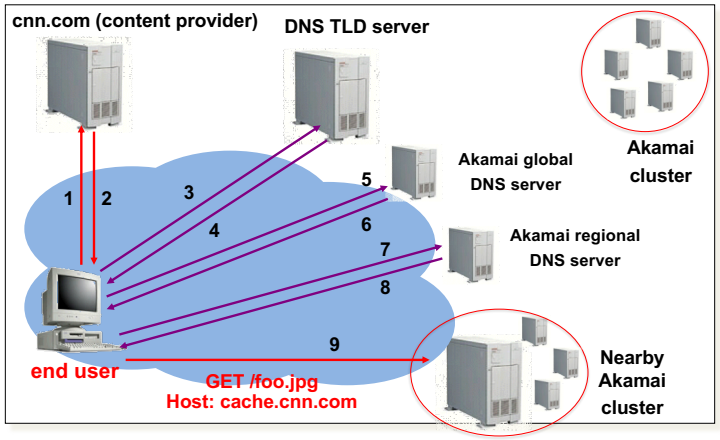


16

# How Akamai Uses DNS

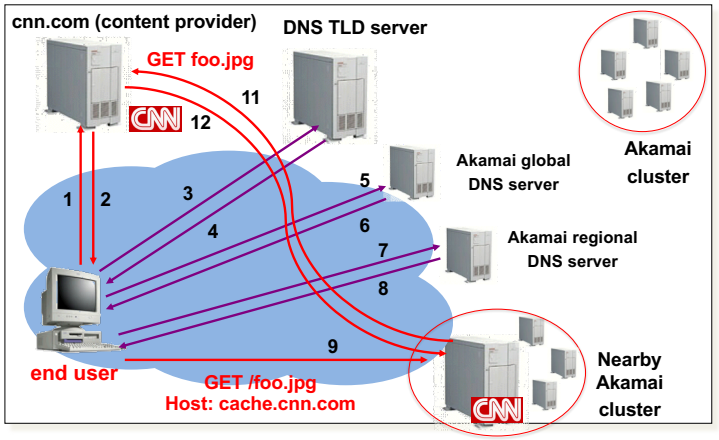

# How Akamai Uses DNS

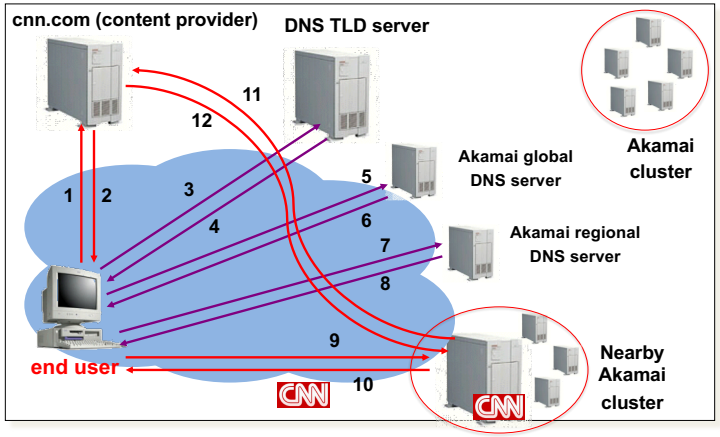




# How Akamai Works: Cache Hit

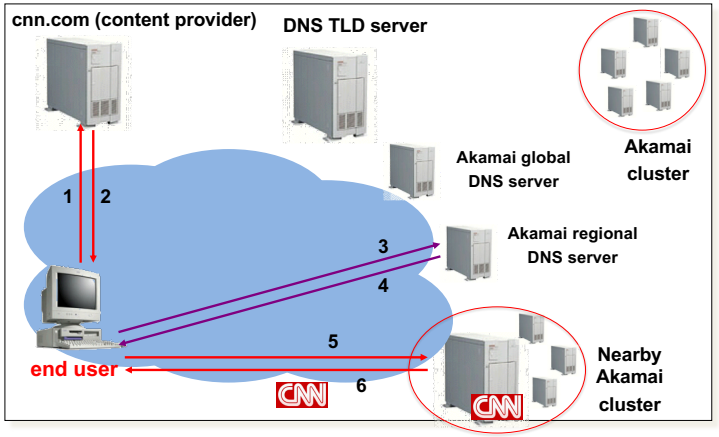


# Mapping System

- Equivalence classes of IP addresses
  - IP addresses experiencing similar performance
  - Quantify how well they connect to each other
- Collect and combine measurements
  - Ping, traceroute, BGP routes, server logs
    - E.g., over 100 TB of logs per days
  - Network latency, loss, and connectivity



# Mapping System

- Map each IP class to a preferred server cluster
  - Based on performance, cluster health, etc.
  - Updated roughly every minute
- Map client request to a server in the cluster
  - Load balancer selects a specific server
  - E.g., to maximize the cache hit rate

## Adapting to Failures

- Failing hard drive on a server
  - Suspends after finishing "in progress" requests
- Failed server
  - Another server takes over for the IP address
  - Low-level map updated quickly
- Failed cluster
  - High-level map updated quickly
- Failed path to customer's origin server
  - Route packets through an intermediate node

21

## Conclusion

- Content distribution is hard
  - Many, diverse, changing objects
  - Clients distributed all over the world
  - Reducing latency is king
- Contribution distribution solutions
  - Reactive caching
  - Proactive content distribution networks

22