Content Distribution Networks (CDNs)

Jennifer Rexford

COS 461: Computer Networks

Lectures: MW 10-10:50am in Architecture N101

http://www.cs.princeton.edu/courses/archive/spr12/cos461/

Second Half of the Course

- Application case studies
  - Content distribution and multimedia streaming
  - Peer-to-peer file sharing and overlay networks
- Network case studies
  - Home, enterprise, and data-center networks
  - Backbone, wireless, and cellular networks
- Network management and security
  - Programmable networks and network security
  - Internet measurement and course wrap-up

Single Server, Poor Performance

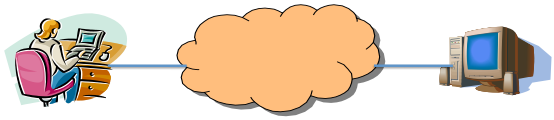

- Single server
  - Single point of failure
  - Easily overloaded
  - Far from most clients
- Popular content
  - Popular site
  - "Flash crowd" (aka "Slashdot effect")
  - Denial of Service attack

Skewed Popularity of Web Traffic

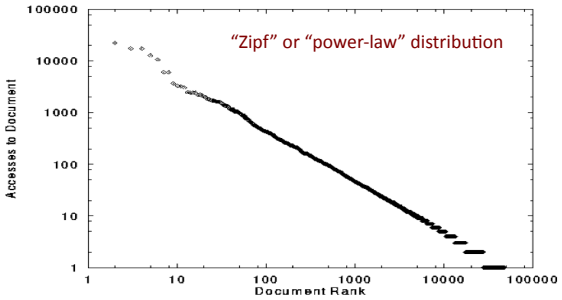

Characteristics of WWW Client-based Traces
Carlos R. Cunha, Azer Bestavros, Mark E. Crovella, BU-CS-95-01

Web Caching

Proxy Caches

- Reactively replicates popular content
- Smaller round-trip times to clients
- Reduces load on origin servers

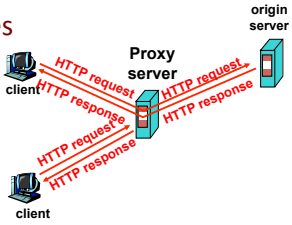

- Reduces network load, and bandwidth costs
- Maintain persistent TCP connections

Forward Proxy

- Cache close to the client
  - Improves client performance
  - Reduces network provider's costs
- Explicit proxy
  - Requires configuring browser
- Implicit proxy
  - Service provider deploys an "on path" proxy
  - … that intercepts and handles Web requests

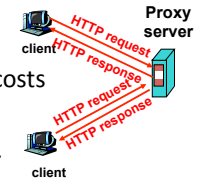

Reverse Proxy

- Cache close to server
  - Improve client performance
  - Reduce content provider cost
  - Load balancing, content assembly, transcoding, etc.
- Directing clients to the proxy
  - Map the site name to the IP address of the proxy

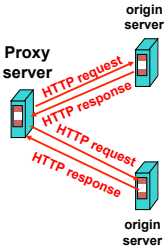

Google Design

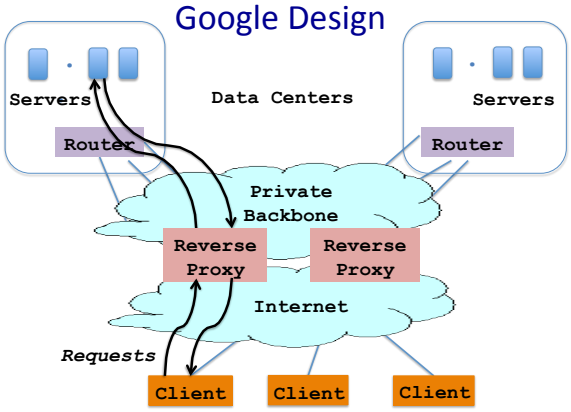

Limitations of Web Caching

- Much content is not cacheable
  - Dynamic data: stock prices, scores, web cams
  - CGI scripts: results depend on parameters
  - Cookies: results may depend on passed data
  - SSL: encrypted data is not cacheable
  - Analytics: owner wants to measure hits
- Stale data
  - Or, overhead of refreshing the cached data

Content Distribution Networks

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

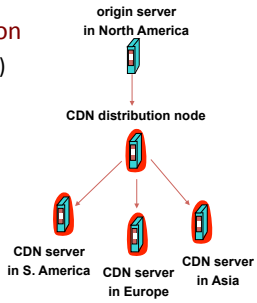

Server Selection Policy

- 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

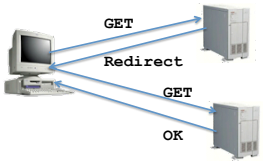

- Advantages
  - Fine-grain control
  - Selection based on client IP address
- Disadvantages
  - Extra round-trips for TCP connection to server
  - Overhead on the server

Server Selection Mechanism

- Routing
  - Anycast routing

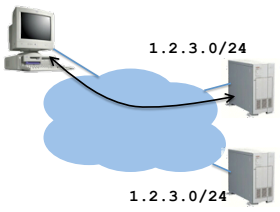

- Advantages
  - No extra round trips
  - Route to nearby server
- Disadvantages
  - Does not consider network or server load
  - Different packets may go to different servers
  - Used only for simple request-response apps

Server Selection Mechanism

- Naming
  - DNS-based server selection

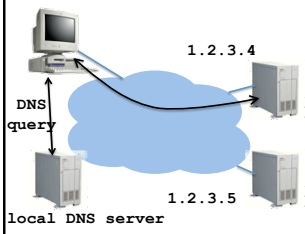

- Advantages
  - Avoid TCP set-up delay
  - DNS caching reduces overhead
  - Relatively fine control
- Disadvantage
  - Based on IP address of local DNS server
  - "Hidden load" effect
  - DNS TTL limits adaptation

How Akamai Works

Akamai Statistics

- Distributed servers
  - Servers: ~61,000
  - Networks: ~1,000
  - Countries: ~70
- Many customers
  - Apple, BBC, FOX, GM IBM, MTV, NASA, NBC, NFL, NPR, Puma, Red Bull, Rutgers, SAP, …
- Client requests
  - Hundreds of billions per day
  - Half in the top 45 networks
  - 15-20% of all Web traffic worldwide

19

How Akamai Uses DNS

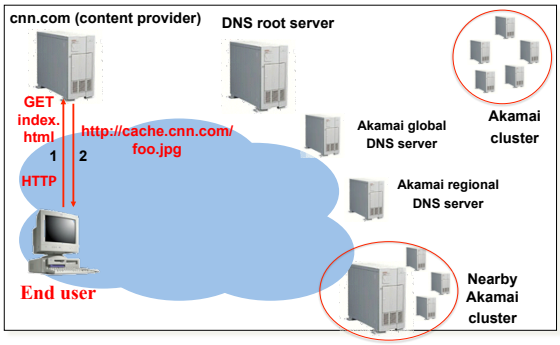

20

How Akamai Uses DNS

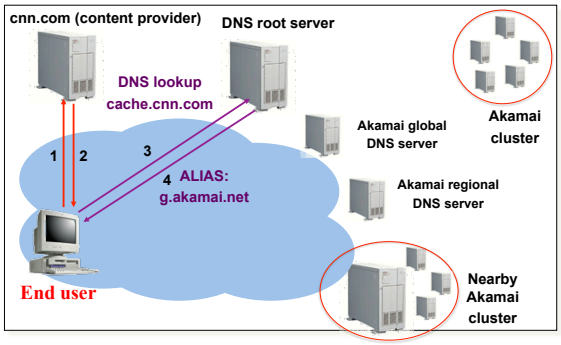

21

How Akamai Uses DNS

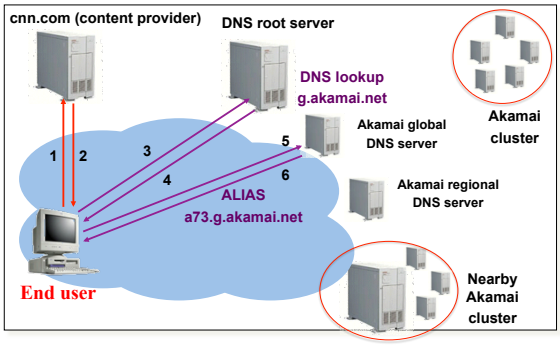

22

How Akamai Uses DNS

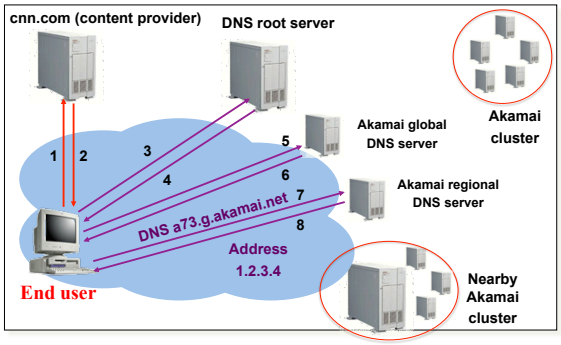

23

How Akamai Uses DNS

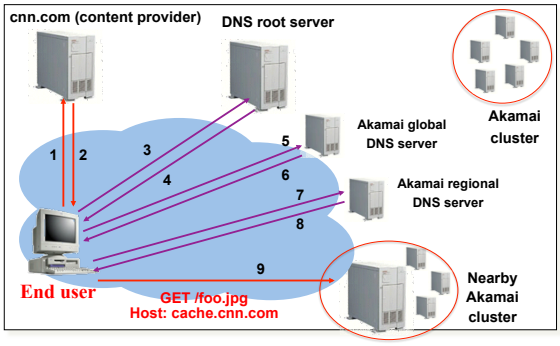

24

How Akamai Uses DNS

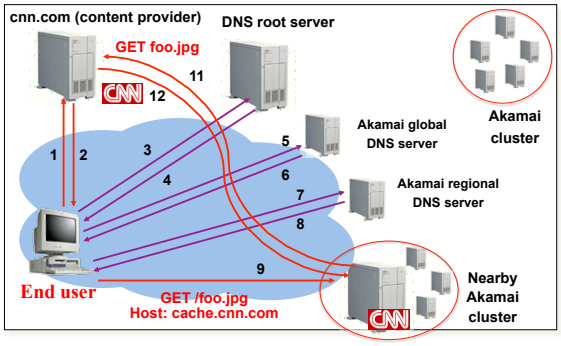

How Akamai Uses DNS

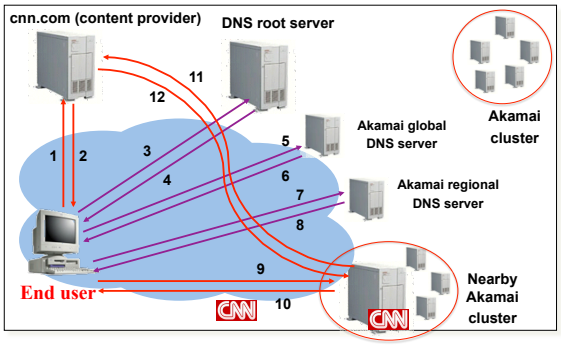

How Akamai Works: Cache Hit

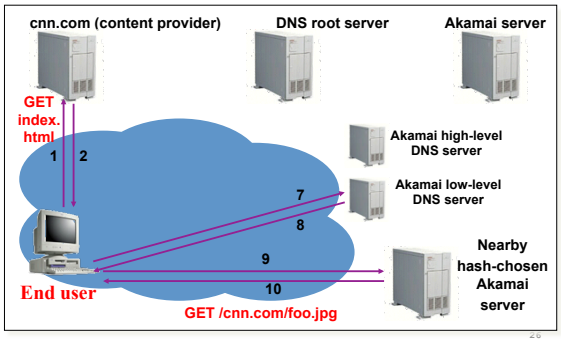

Mapping System

- 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

Akamai Transport Optimizations

- Bad Internet routes
  - Overlay routing through an intermediate server
- Packet loss
  - Sending redundant data over multiple paths
- TCP connection set-up/teardown
  - Pools of persistent connections
- TCP congestion window and round-trip time
  - Estimates based on network latency measurements

Akamai Application Optimizations

- Slow download of embedded objects
  - Prefetch when HTML page is requested
- Large objects
  - Content compression
- Slow applications
  - Moving applications to edge servers
  - E.g., content aggregation and transformation
  - E.g., static databases (e.g., product catalogs)
  - E.g. batching and validating input on Web forms

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
- Next time
  - Multimedia streaming applications