

Distributed Systems

16. Content Delivery Networks (CDN)

Paul Krzyzanowski

pxk@cs.rutgers.edu

Motivation

- Serving web content from one location presents problems
 - Scalability
 - Reliability
 - Performance
- “Flash crowd” problem
 - What if everyone comes to you at once?
- Cache content and serve requests from multiple servers at the network edge (close to the user)
 - Reduce demand on site’s infrastructure
 - Provide faster service to users
 - Content comes from nearby servers



Akamai Distributed Caching

- Akamai runs on 95,000 servers in 1,900 networks across 71 countries

Goal

- Company evolved from MIT research
- "Invent a better way to deliver Internet content"
- Tackle the "flash crowd" problem

Traditional approaches

- Local clustering
 - Data center or ISP can fail
- Multihoming
 - IP protocols (BGP) are often not quick to find new routes
- Mirroring at multiple sites
 - Synchronization can be difficult
- Proxy servers
 - Typically a client-side solution
 - Low cache hit rates
- All require extra capacity and extra capital costs

Akamai's goal

- Akamai tries to serve clients from servers likely to have the content
 - Nearest: lowest round-trip time
 - Available: server that is not too loaded
 - Likely: server that is likely to have the data

Akamai Mapping

- Akamai does this via Dynamic DNS
- Direct requests to the right content servers
- Resolve a host name based on:
 - service requested (e.g., QuickTime, HTML, Windows Media)
 - content requested (is server likely to have it? – based on hash)
 - server health
 - server load
 - user location
 - network status
 - load balancing

Akamai needs data to do this

- Map network topology
 - Based on BGP and traceroute information
 - Estimate hops and transit time
- Content servers report their load to a monitoring application
- Monitoring app publishes load reports to a local DNS server
- DNS server determines which IP addresses to return when resolving names
- Load shedding:
 - If servers get too loaded, the DNS server will not respond with those addresses

DNS Resolution

- Resolve a7.g.akamai.net
- NS resolver contacts root server
 - Root server sends a referral to a name server responsible for .net
- Resolver queries .net name server
 - Returns a referral for .akamai.net
 - This is the top-level Akamai server
- Resolver queries a top-level Akamai server
 - Returns a referral for .g.akamai.net
 - Low-level Akamai server (TTL approx 1 hour)
 - Low-level servers are in the same location as edge servers closest to user
- Resolver queries a low-level Akamai server
 - Returns IP addresses of servers available to satisfy the request
 - Short TTL (several seconds to 1 minute)

Populating the cache

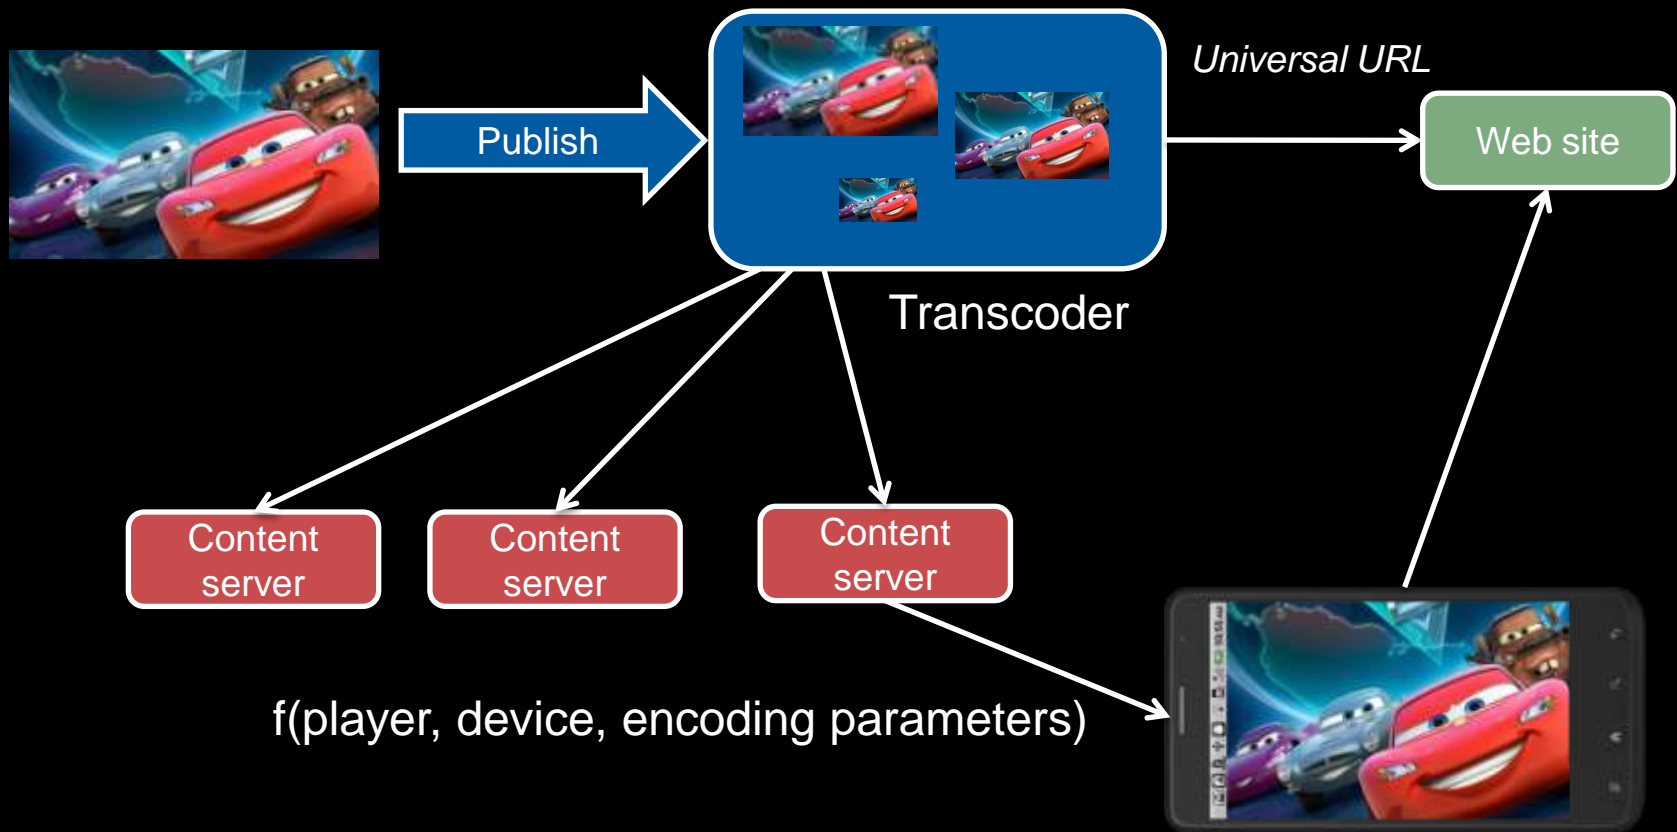
- If a server does not have the content
 - Get it from another Akamai server
 - Or get it from the site

Types of content

- **Static content**
 - Cached depending on original site's requirements (never to forever)
- **Dynamic content**
 - Caching proxies cannot do this
 - Akamai uses Edge Side Includes technology (www.esi.org)
 - Assembles dynamic content on edge servers
 - Similar to server-side includes
 - Page is broken into fragments with independent caching properties
 - Assembled on demand
- **Streaming media**
 - Live stream is sent to an entry-point server in the Akamai network
 - Stream is delivered from the entry-point server to multiple edge servers
 - Edge servers serve content to end users.

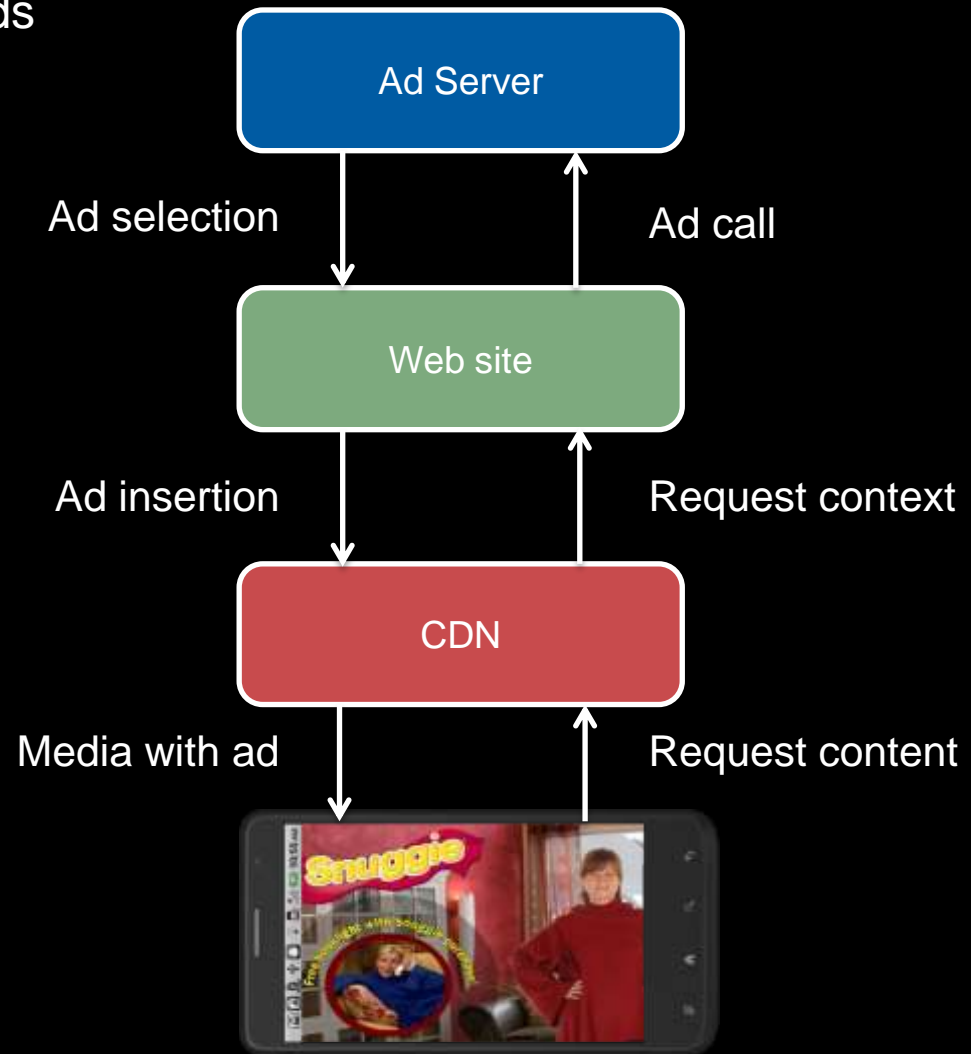
Other Things CDNs Do

LimeLight Reach Video



Server-side Video Ad Insertion

Example, Limelight Reach Ads



The End