

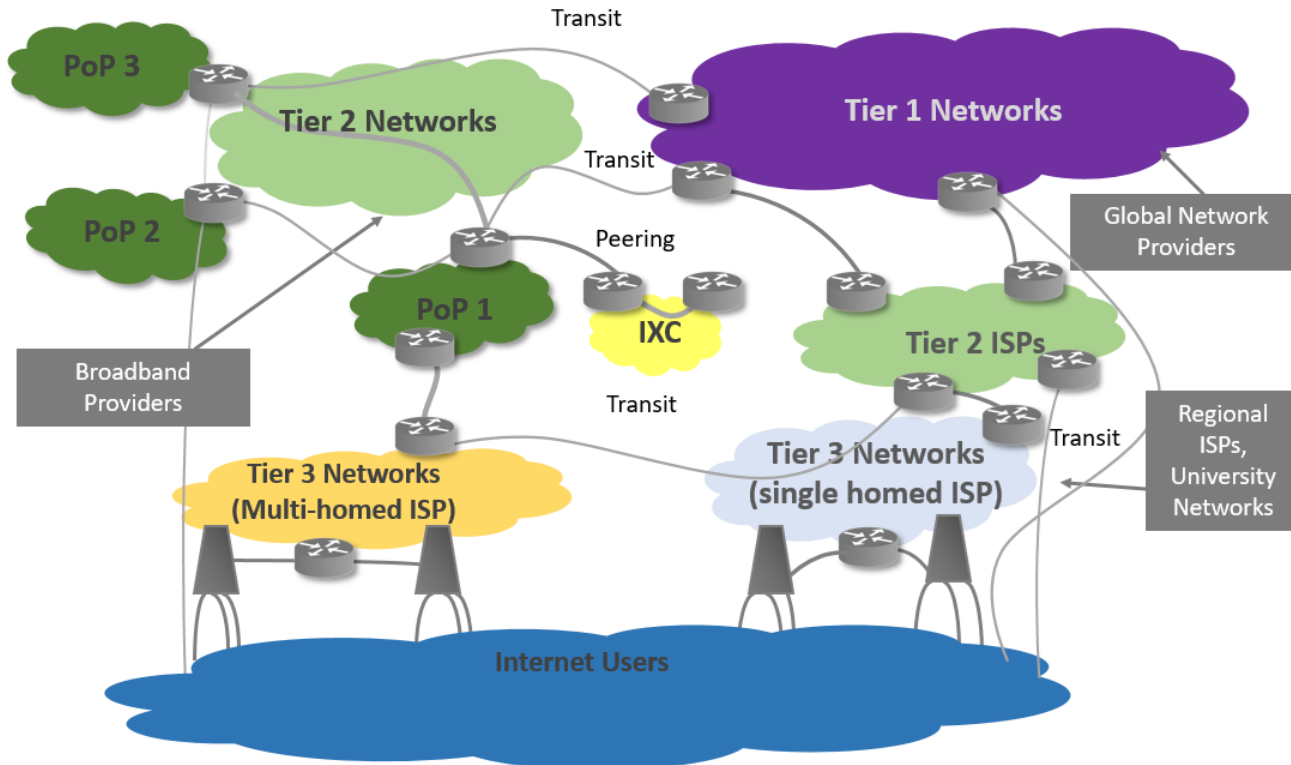


Content delivery networks

Antoni Bertel



Internet is complex



Challenges

Distance

Network bottlenecks

Packet loss

Congestion

TCP slow start

Round-trip delay time

Retransmission

Chatty protocol

CDN - content Delivery Network

Content Delivery Networks (CDN), which evolved first in 1998, replicate content over several mirrored (surrogate) Web servers strategically placed at various locations in order to deal with the *flash crowds*. Geographically distributing the Web servers' facilities is a method commonly used by service providers to improve performance and scalability.

Flash crowds is a 1973 English novella. In network is a sudden, large surge in traffic to a particular WEB resource.

CDN - content Delivery Network

1st Generation: Focused on Static of Dynamic Web Documents

2nd Generation: Focused on Video-on-Demand, audio and video streaming.

CDN - Why it is done?

Accelerated Web Performance - CDNs help to improve the delivery of website content (static and dynamic) so that websites perform better, load faster, and generate more revenue for content providers.

Software Updates and Downloads - CDNs enhance the automatic or on-demand delivery of software or file downloads, including software patches.

Rich Media Content Streaming - CDNs help deliver rich media content, that is, interactive digital media such as audio and video files in different encoding formats (HTTP adaptive streaming, Adobe HDS, and Apple HLS, etc.), to specialized streaming clients and devices of end-users.

CDN - Why it is done?

IPTV Use Case - CDNs are used for Internet Pay TV (IPTV), such as catch-up TV service for recent programs of broadcasters, as well as online live video channels.

Managing and Delivering User-Generated Content - CDNs are often used to facilitate the hosting and distribution of user-generated content. YouTube is the prime example where a CDN is used to deliver content uploaded by end-users.

End-to-End Online Video - CDNs at its fullest use are often employed as the glue in integrating online video and content management platform to create an end-to-end delivery chain, offering a real differentiation through high quality user experience.

Target Audience

Networking companies

Data center providers

Cloud Service Providers (CSPs)

Internet Service Providers (ISPs)

Telecom providers

Value Added Resellers (VARs)

Managed Service Providers (MSPs)

Hosting vendors

Market segments

E-commerce and Consumer Products/Retail - CDNs offer whole-site and dynamic content delivery services to website owners. A large portion of such websites is composed of applications and dynamic content, such as e-commerce/online retailers (eBay, Amazon), auction website (eBay), and consumer products website

Media and Entertainment - . CDNs offer progressive download, linear, VoD, and live streaming services to digital media companies such as content broadcasters (CNN, BBC), Internet-based publishing (Fairfax Media), and experimental digital media.


Market segments

Hospitality, Travel, and Leisure - This market segment features websites and services used by end-users for accommodation booking, holiday booking, air ticket purchase, and so on.

Banking, Financial Services, and Insurance - CDNs help this market segment by delivering a large amount of static website content, as well as dynamic content and applications, such as currency exchange, financial projections, and loan calculators.

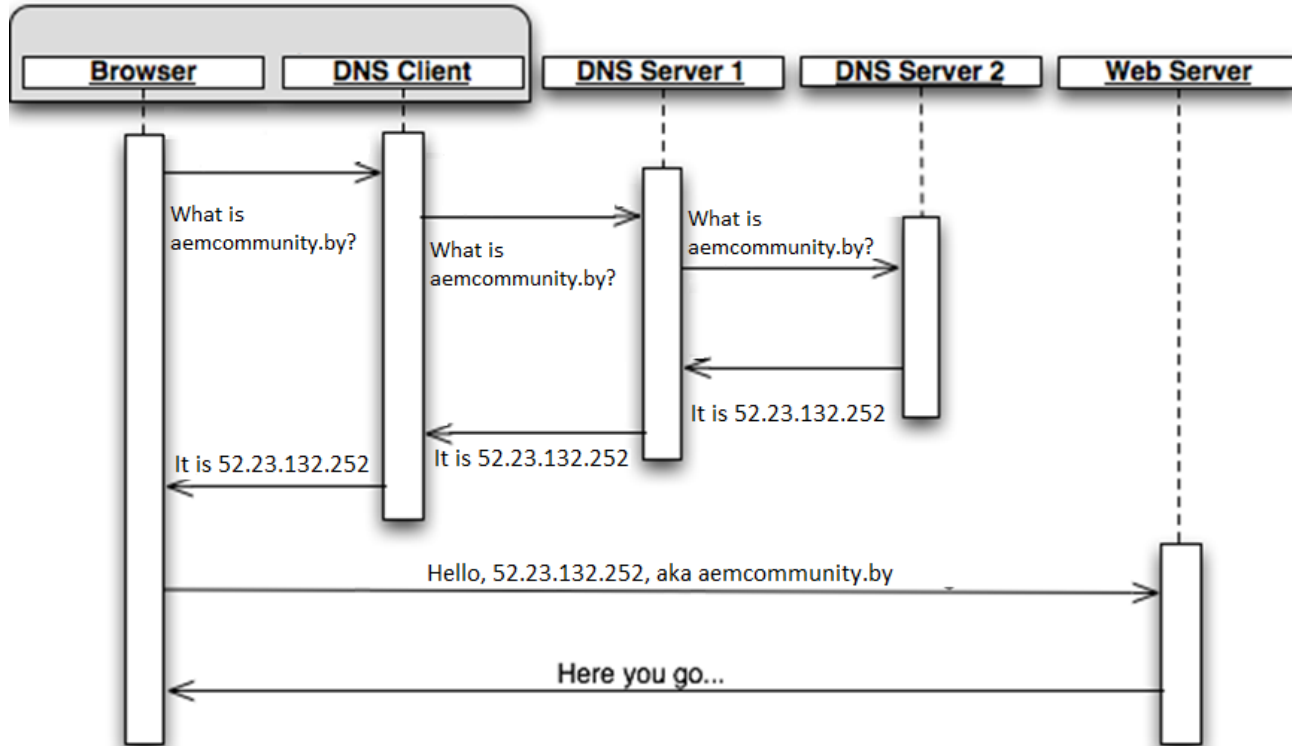
Public Sector/Government and Enterprise - CDNs also help the Government (such as the president.gov.by) to deliver heavy websites of secure and static content.

CDN architecture




| | |
|----------------|---|
| Publisher | A provider of Data to Consumers that uses the services of the Network |
| Authorize | the Publisher gives the Network permission to deliver particular Data |
| Report | the Publisher requests statistics on the services provided by the Network |
| Origin | the Publisher provides an original copy of the Data to the Network |
| Data (Content) | A piece of digital information that the Publisher has created or licensed for distribution to Consumers |
| Request | the Publisher directs the Consumer to request the Data from the Network |
| Deliver | the Network provides the Consumer with the Data/Content |
| Subscriber | User of the Publisher's Data/Content |

DNS-normal routing



DNS-based request routing




Once a DNS query for an outsourced URL arrives at the CDN's ADNS,

The CDN provider can direct different HTTP requests to different content servers,

by returning different IP addresses to different queries.

Request Routing Using IP Anycast



The same IP address is assigned to multiple endpoints,

and packets to this IP address are routed to the endpoint that is the nearest to the sender based on the shortest (or otherwise most preferred) path from the network-routing perspective

CDN overlays

Increases the security and mitigates distributed denial of service (DDoS) attacks on websites and other online services

**Security
overlay**


Provides wide-area communication with more reliability, lesser latency, and greater throughput than the public Internet can

**Routing
overlay**

Deliver websites, on-demand videos, music downloads, software downloads, and other forms of online content

**Caching
overlay**

Performance benefits of a caching overlay



Analytic overlay

Know It All - Access constantly updated information about server performance, user demographics, and bandwidth utilization

Know It Now - With live events, often there is no time to wait for server log processing. Use real-time statistics and take action—in the moment.

Understand the Audience - See how users are interacting with the website content, where they are coming from, and how long they stick around

What exactly should be distributed?

The user behavior when requesting data and services on the Internet is following Zipf-like distributions, which concentrate a major portion of activity on a small subset of the most popular items. Such behavior makes it efficient to distribute popular data over caches and servers even of limited size and to deliver data from many sources that are located closer to the user.

$$A(R) = aR^{-\beta} (a > 0; \beta > 0)$$

$\alpha = A(1)$ is the maximum number of requests observed for an item in the statistic, $A(R)$ is the number of requests to an item, R - rank of item, β - exponent, calculated using least squares method

Cache replacement Strategies

First In First Out (FIFO)

Last In First Out (LIFO)

Least Recently Used (LRU)

Most Recently Used (MRU)

Pseudo-LRU (PLRU)

Random Replacement (RR)

Segmented LRU (SLRU)

LFU with Dynamic Aging (LFUDA)

Low Inter-reference Recency Set (LIRS)

Adaptive Replacement Cache (ARC)

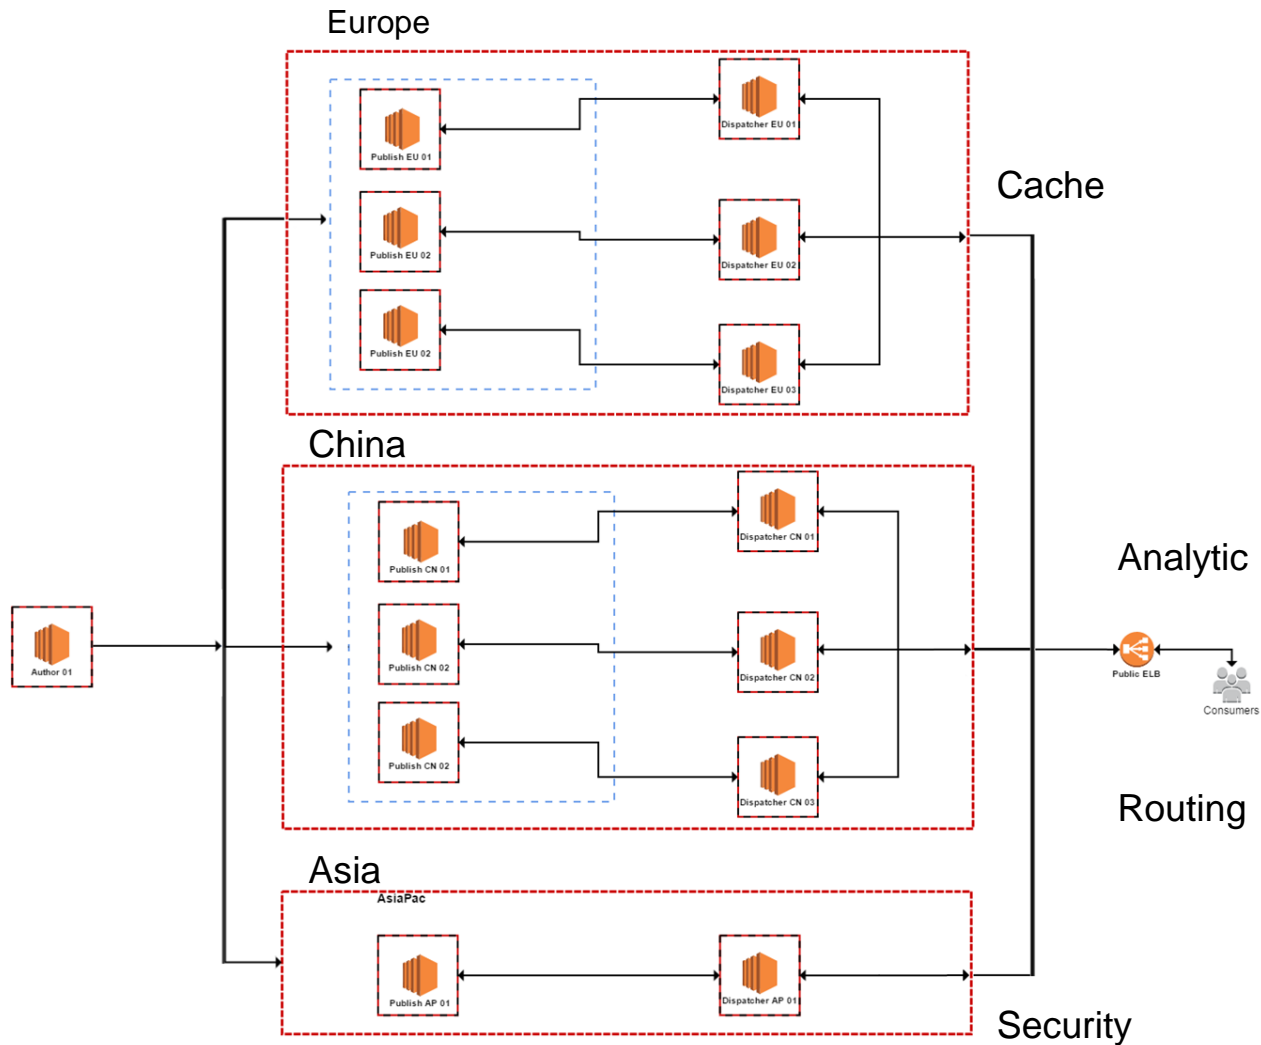
Clock with Adaptive Replacement (CAR)

Multi Queue (MQ) caching
algorithm | Multi Queue (MQ)

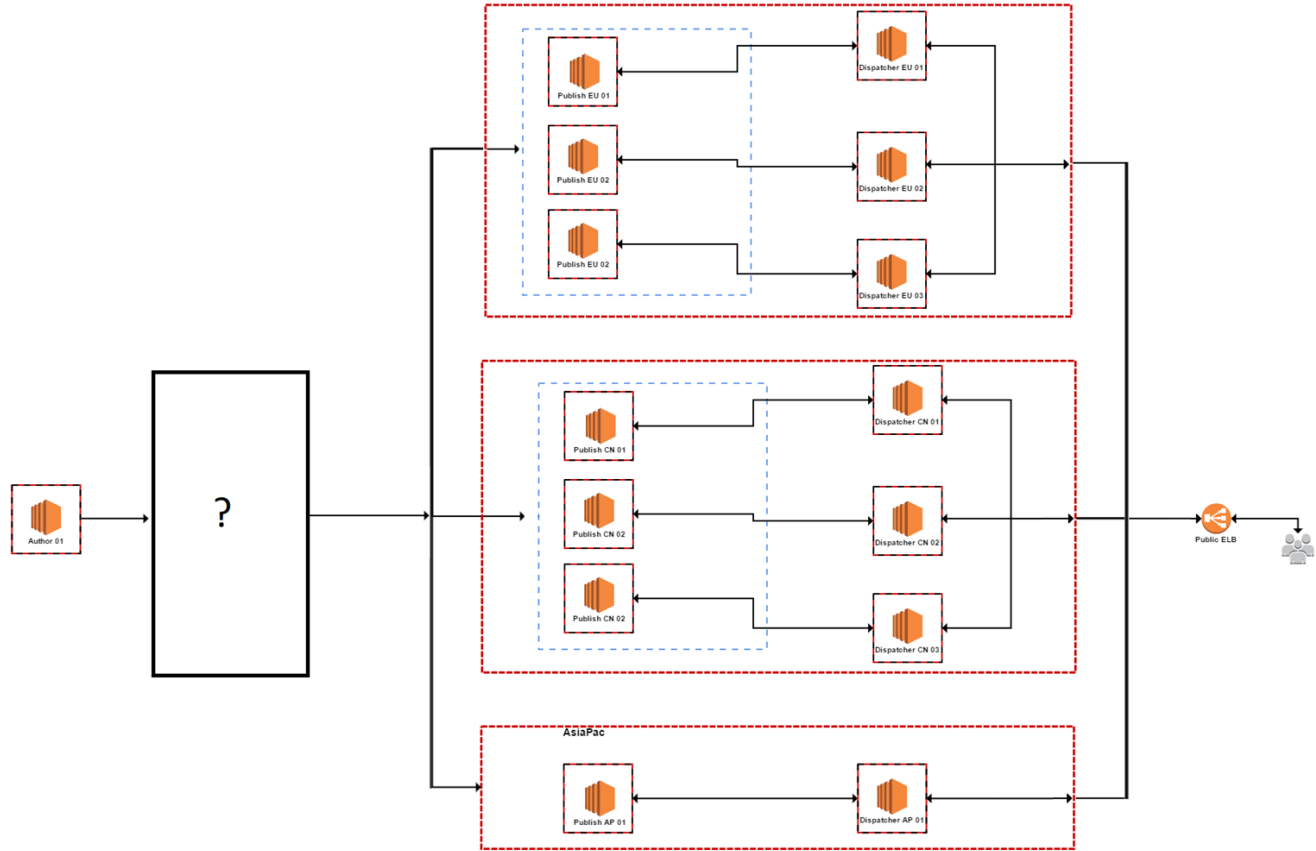
Custom CDN

Issues

1. Network bottlenecks
1. Heavy scaling
1. Losing money
1. Long-Running Request-Response



Solution





See in action





See in action



Cloud solitons



Care about

Uptime/availability of service

Throughput performance

First byte delivery time

Reporting/analytics

Range of products (Analytic, Dynamic Tag Management, etc.)

Customer service



Question Time

