

Embedded CDNs in 2023

-A status

Nina Bargisen
Steve Meuse
September 2023

A quick introduction

- This presentation was largely produced by by colleague, Nina Bargisen.
- My first time back to Australia since 2000 IETF in Adelaide!
- Today's topic will be largely about Embedded CDN caches in Service Provider networks


but a quick story first!

The early days 1999-2000

- In 1999, CDNs were a new concept, Akamai was the main player.
- Content was still mostly delivered from Hosting Providers (and on-prem)
- I was working for GTE Internetworking, which ran Internet AS1
- We had a small POP in Sydney to support some large, multinational, companies (AS201)
- This POP was server by 2xE1!!! (and had several downstream customers)
- An effort was made to conserve bandwidth.....

...so we used a different kind of cache!

Transparent Web Caches!



- 1x100Mb/s interfaces
- Ran a predecessor to wccp
- Infolibra
- Inktomi
- 9600 baud modem
- 10Mb/s management network
- We tested 3-4 vendors and opensource to see which worked best

Did it work?

NO!

**Our (smart) customers were already
caching!**

What is a CDN?

- ✓ A collection of geographically distributed servers
- ✓ A method to place content on the servers
- ✓ A method to steer end-users to the closest servers

Embedded servers are when some of the CDNs servers are located within ISPs networks and intended to only serve end-users in that network. The servers will use address space originated in the ISPs network and not in the CDNs ASN(s).

Clustering is when a number of servers share their storage. This means that the content is distributed over the servers in the cluster and will not be present at all of them. Sometimes servers in different sites can be clustered together.

A bit about the CDN market

Growth in the market based on the increased consumer demand for entertainment and OTT services and in the later years on increased demand for e-learning and work from home solutions.

The traditional players like Akamai, Lumen, Tata, Edgio, etc are well established as reliable and with global reach.

The newer major CDNs like Cloudflare, Fastly and Stackpath are offering a more diverse and specialized set of services like security services and DDOS protection, high performance real time delivery services etc.

Which is also offered by Akamai

Large content producers who use commercial CDNs often use multiple CDNs. This both to support their different types of content, to leverage different geographical scope of the CDN, create resilience and finally - more than one provider means better leverage when negotiating contracts...

A bit about the CDN market

The giant content producers like Netflix, Apple, Microsoft, Amazon, Facebook and Google have all build their own specialised CDN to support their core services (and some also sell CDN services...)

Building a CDNs on a global scale is not small task, so this solution is only for the largest and most resourceful content producers with enough demand for traffic to make the business case work.

The cloud providers almost all have a CDN to front their content

The CDNs are almost all also offering compute and other cloud services

So in the end they end up offering the same services but with different specialities and strengths

Who offers embedded servers?

- Akamai
- Netflix
- Google
- Amazon
- Facebook
- Cloudflare
- CDN77
- Microsoft
- Apple
- Qwilt
- And more to come

Statistics

- Petros Gigis, Lefteris Mannasakis and more presented at Sigcomm'21 that over 3,500 ASNs globally have embedded servers
- Most ISPs in the Kentik traffic sample have 6 or less embedded CDNs
- The **top 6** in the Kentik sample are:

Google (25.5%)


Netflix (26.5%)

Facebook (16.2%)

Akamai (11.6%)

Apple (3.3%)

CDN77 (3.3%)



Percentages of the ISPs in the sample that have embedded servers from the CDN

A note on localization...

Latin America

In our observations, **Latin America** has a higher number of unique CDN cache providers, why?

Kentik has observed as many as 10 unique cache providers in LATAM.

- ✓ Higher cost to deliver bits?
- ✓ Fewer legacy video transport services?
- ✓ Logistical challenges?

A note on localization...





ANZ Region

Kentik has observed that the **ANZ region** is much more similar to the North American region with only the main players

The average is 4 unique providers (Netflix, Akamai, Google, Facebook), with a few others like Apple and Valve.

- ✓ Differences in Interconnect policy?
- ✓ Population Density?
- ✓ Fewer Logistical challenges?

Total distribution of traffic marked as CDN traffic


-  ~74% over the external network border
-  ~14% from embedded servers
-  ~7% embedded to embedded internally in the networks
-  ~5% Externally to embedded



Average traffic over 24 hours that includes a Sunday evening.


Statistics

Google



- Outside to end-users (39%)
- Embedded to end-users (56%)
- Outside to embedded (5%)
- Embedded to embedded (1%)

Netflix



- Outside to end-users (22%)
- Embedded to end-users (72%)
- Outside to embedded (3%)
- Embedded to embedded (3%)

Facebook



- Outside to end-users (43%)
- Embedded to end-users (54%)
- Outside to embedded (4%)
- Embedded to embedded (0%)

Akamai



- Outside to end-users (73%)
- Embedded to end-users (18%)
- Outside to embedded (9%)
- Embedded to embedded (0%)

Steer to the closest cache

BGP

- BGP is used to signal the subnets that a cache or cluster should serve
- The most control is given for the systems where every cache has a BGP session and the steering system respects common BGP parameters.
- A setup where all Servers have an identical session will work for most deployments.
- Prefixes are sent from the ISP to the cache

DNS Server

- End-users are mapped to a cluster based on the DNS server they are using to request the content

Anycast

- Anycast addresses are announced from the caches and the ISPs routing decides which is the closest

Geolocation

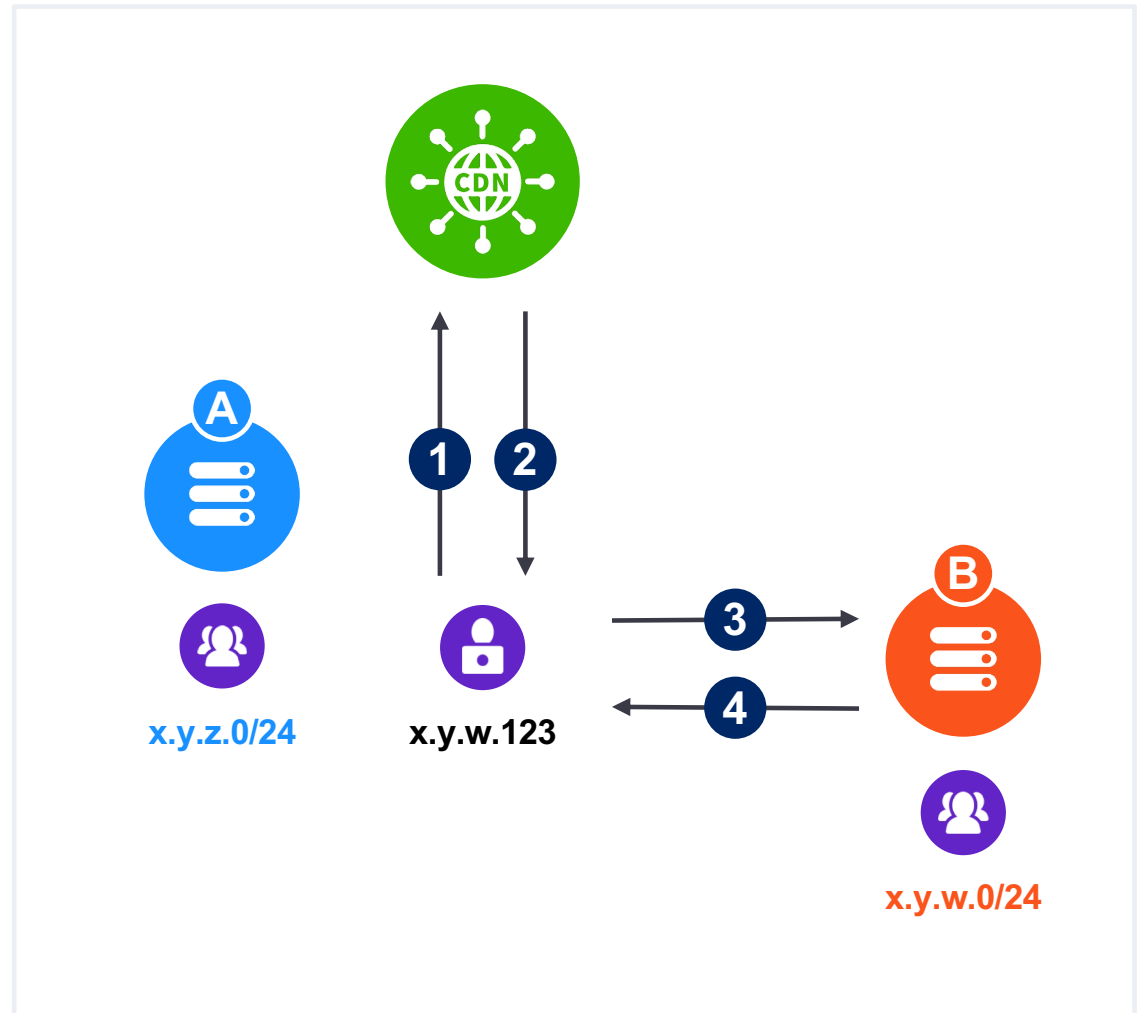
- Geolocation is used by some as the primary way of mapping an end-user to a cluster. Some use the geolocation of the end-user IP address, but some rely on the geo location of the DNS server the end-user is using to request the content.

BGP

- **x.y.z.0/24** is announced to **A**
- **x.y.w.0/24** is announced to **B**

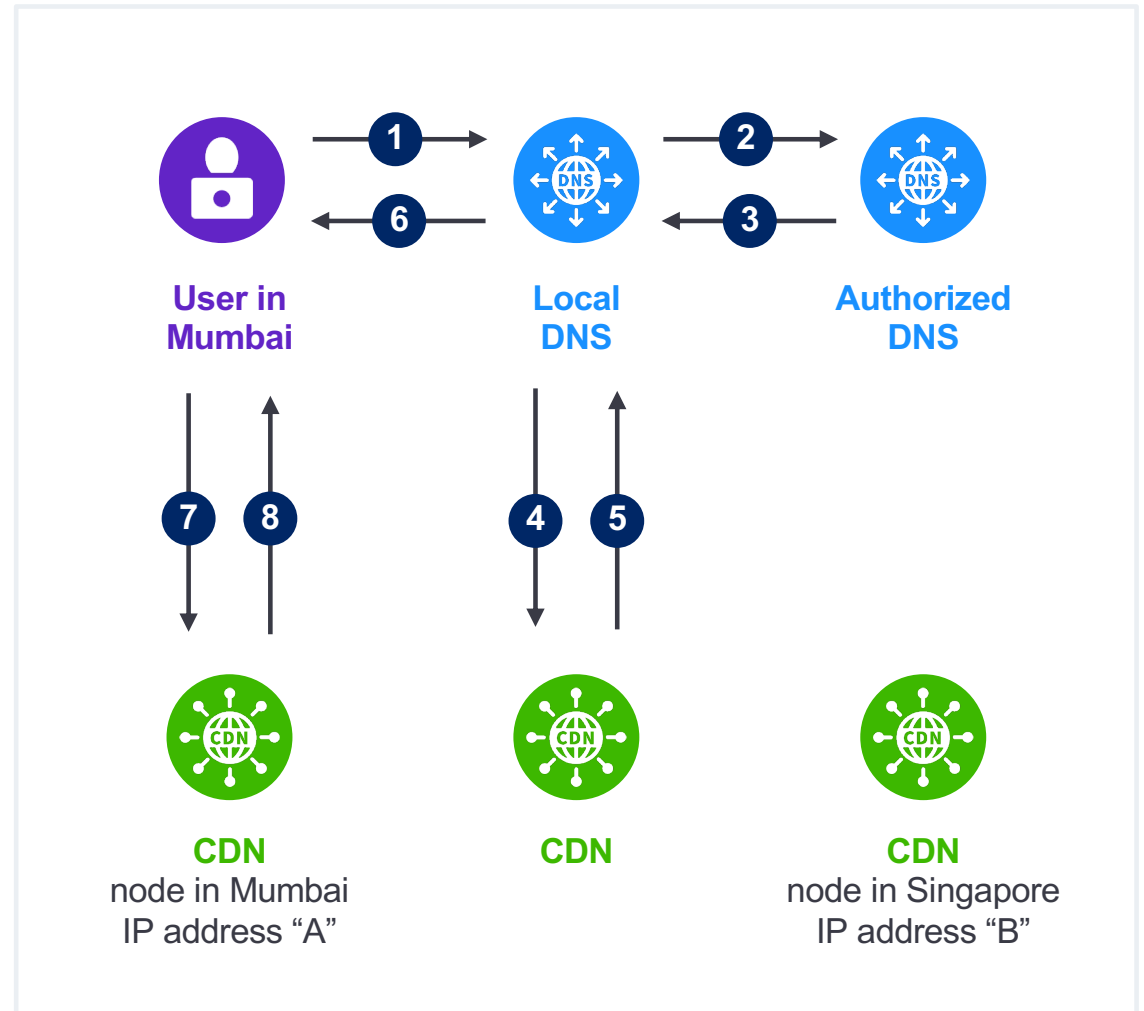
1. Give me movie
2. Go to **B** and get movie
3. Give me movie
4. Movie

This works really well for ISPs who have a very strict regional address plan. It does not work for ISPs where they do not have control over IP address assignments or who do not have regional aggregation



DNS

1. Where is **site.com**?
2. Where is **site.com**?
3. **site.com** is **site.com.cdnsomething.com**
4. Where is **site.com.cdnsomething.com**
5. **site.com.cdnsomething.com** is **A** for you
6. **Site.com** is **A**
7. Give me **Site.com**
8. **Site.com**



And then some of the magic

Latency and other QoS data

- The CDN creates a mapping of ISP DNS servers to POPs or clusters based on latency measurements
- Some CDNs include internet connectivity from the cluster to the DNS servers or IP addresses in their mapping algorithm
- The content owned CDNs also include signals from the content clients in the algorithm

Load


- The load of the individual servers or clusters is part of the decision of exactly which server the end-user will end up using.
- Load balancing can take place in the cluster or in the CDN response to (where is?)

Content

- Not all content is on all servers
- Content that is pushed out is distributed according to CDN magic
- Not all content from a CDN will be cached when the content is pulled
- This means where the content is/can be is also part of the steering decision.

Traffic profiles for different ways of placing content


- Centrally calculated placement
- Fill window
- Optimized to reduce edge fill traffic as much as possible



Typical Netflix traffic profile in well dimensioned deployment

Traffic profiles for different ways of placing content

- Proxy and Akamai magic
- Relatively dynamic content
- Not all content will be at the embedded servers



Typical Akamai traffic profile in well dimensioned deployment


Traffic profiles for different ways of placing content

- Proxy and Facebook magic
- Relatively dynamic content
 - *Only the static is cached*
- Not all content will ever be at the embedded servers

Feeds

Likes

Anything that needs to be globally consistent



Typical Facebook traffic profile in well dimensioned deployment

Limit who can use the Servers



Most embedded Servers will have access list that limits who will be served by the cache



The common source to generate these is a BGP feed from the host ISP



Some will use one single feed per ISP, other per cluster or per cache.

There must be consistency between the user-group that can be mapped to a cluster or server and the user-group that are allowed to use the cluster or server.

This means there is a requirement for the CDNs where each server has a BGP sessions that the announcements to each server in a cluster is identical.

Deploying embedded servers deep into the network

- Understand if the end-user mapping is supported by your network implementation
 - DNS servers
 - IP address plan
 - Use of CGN
- Be aware of clustering
- Off load/hit rate for the local cluster might be lower than for a big cluster because of less disk space in the cluster



North American Market

Meanwhile, in North America.....

MEDIA

Amazon's 'Thursday Night Football' begins a new era of streaming for the NFL

PUBLISHED FRI, AUG 19 2022-7:00 AM EDT | UPDATED SAT, AUG 20 2022-4:15 PM EDT



Alex Sherman
@SHERMAN4949

SHARE [f](#) [t](#) [in](#) [✉](#)

Media & Entertainment


Amazon Prime Video's 'Thursday Night Football' starts strong with 15.3 million viewers

Lauren Forristal @laurenforristal / 1:46 PM EDT • September 22, 2022


[Comment](#)



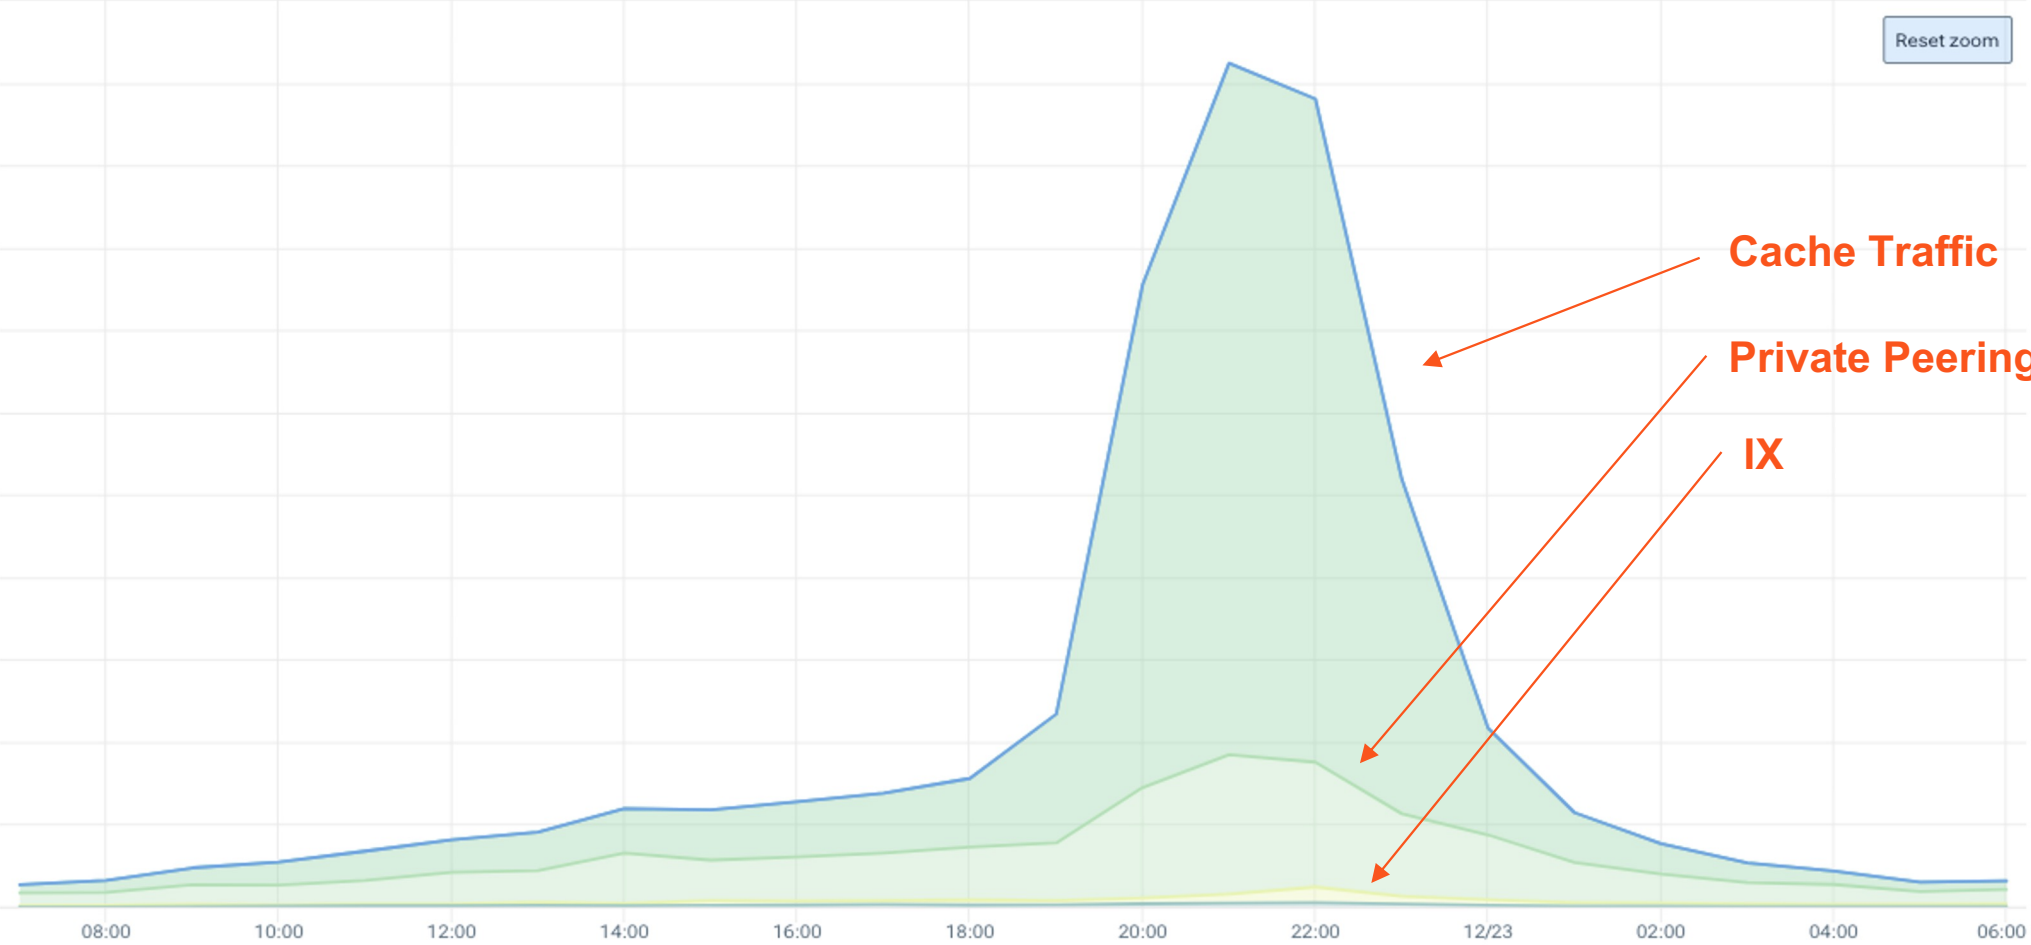
Amazon Prime Video - Embedded Cache traffic only




Amazon Prime Video - Embedded Cache + Peering/Transit



Amazon Prime Video - Embedded Cache + Peering/Transit



Amazon Prime Video - External Traffic - CDN breakdown



Reset zoom

Amazon
Cloudfront
Edgio
Akamai
Fastly

And what about the hardware?

- ✓ 1-2U
- ✓ 10G/50G/100G Interfaces
- ✓ Choice of AC and DC power
- ✓ It varies whether a CDN managed switch is required optional or not applicable.
- ✓ Installation process are more/or less automated, but some still require the ISP tech to do basic configuration of the CDN equipment.

Evolves continuously: less space, more throughput, more and faster storage Interfaces supports the industry standard at any given time

CONCLUSION

Is it easier or harder in 2023 compared to 2012?



2012

- Large boxes
- Offload unpredictable
- Few players
- No CGN support

2023

- Small boxes
- More predictable offload
- More predictable and flexible end-user mapping
- Many players
- Some CGN support



What is *your* experience?

Come find us at our booth



Thank You!

Steve Meuse

Service Provider Architect

steve@kentik.com