



Peering with Content Distribution Networks

Patrick W. Gilmore
UKNOF
May 21, 2008

Agenda



- Disclaimer
- What is a Content Distribution Network
- Why CDNs peer with ISPs
- Why ISPs peer with CDNs
- Peering at Multiple Locations
- Questions

Disclaimer



I work for Akamai Technologies, a Content Distribution Network

While all of the information here is applicable to at least one CDN, I cannot guarantee it is applicable to every CDN

In Other Words: Your Mileage May Vary

Feel free to ask me questions about specific CDNs after the talk

What is a Content Distribution Network?



- The RFCs and Internet Drafts define a Content Distribution Network, "CDN", as:

Content Delivery Network or Content Distribution Network. A type of CONTENT NETWORK in which the CONTENT NETWORK ELEMENTS are arranged for more effective delivery of CONTENT to CLIENTS.

What is a Content Distribution Network - In English?



- A CDN is an overlay network, designed to deliver content from the optimal location
 - In many cases, optimal does not mean geographically closest
- CDNs are made of distinct, geographically disparate groups of servers, with each group able to serve all content on the CDN
 - Servers may be separated by type
 - E.g. One group may serve Windows Streaming Media, another group may serve HTTP
 - Servers are not typically shared between media types

What is a Content Distribution Network - In English?



- Some CDNs are network owned (Level 3, Limelight, at&t), some are not (Akamai, Mirror Image, CacheFly, Panther Express)
- Network owned CDNs have all / most of their servers in their own ASN
- Non-Network CDNs place servers directly in other ASNs
 - This means things like NetFlow will not be useful for determining traffic to/from non-network CDNs

How CDNs Work



- When content is requested from a CDN, the user is directed to the optimal server
 - This is usually done through the DNS, especially for non-network CDNs
 - It can be done though anycasting for network owned CDNs
- Users who query DNS-based CDNs be returned different A records for the same hostname
- This is called "mapping"
- The better the mapping, the better the CDN

How CDNs Work



- Example of CDN mapping
 - Notice the different A records for different locations:

```
[NYC]% host www.symantec.com
```

```
www.symantec.com    CNAME    a568.d.akamai.net
a568.d.akamai.net  A        207.40.194.46
a568.d.akamai.net  A        207.40.194.49
```

```
[Boston]% host www.symantec.com
```

```
www.symantec.com    CNAME    a568.d.akamai.net
a568.d.akamai.net  A        81.23.243.152
a568.d.akamai.net  A        81.23.243.145
```


How CDNs Work



- CDNs use multiple criteria to choose the optimal server
 - These include standard network metrics:
 - Latency
 - Throughput
 - Packet loss
 - These also include things like CPU load on the server, HD space, network utilization, etc.
- Geography still counts
 - That whole speed-of-light thing
 - Should be able to solve that with the next version of ethernet...

Why CDNs Peer with ISPs



- The first and foremost reason to peer is improved performance
 - Since a CDN tries to serve content as “close” to the end user as possible, peering directly with networks (over non-congested links) obviously helps
- Peering gives better throughput
 - Removing intermediate AS hops seems to give higher peak traffic for same demand profile
 - Might be due to lower latency opening TCP windows faster
 - Might be due to lower packet loss

Why CDNs Peer with ISPs



- Redundancy
 - Having more possible vectors to deliver content increases reliability
- Burstability
 - During large events, having direct connectivity to multiple networks allows for higher burstability than a single connection to a transit provider
- Burstability is important to CDNs
 - One of the reasons customers use CDNs is for burstability

Why CDNs Peer with ISPs



- Peering reduces costs
 - Reduces transit bill (duh)
- Network Intelligence
 - Receiving BGP directly from multiple ASes helps CDNs map the Internet
- Backup for on-net servers
 - If there are servers on-net, the IX can act as a backup during downtime and overflow
 - Allows serving different content types

Why ISPs peer with CDNs



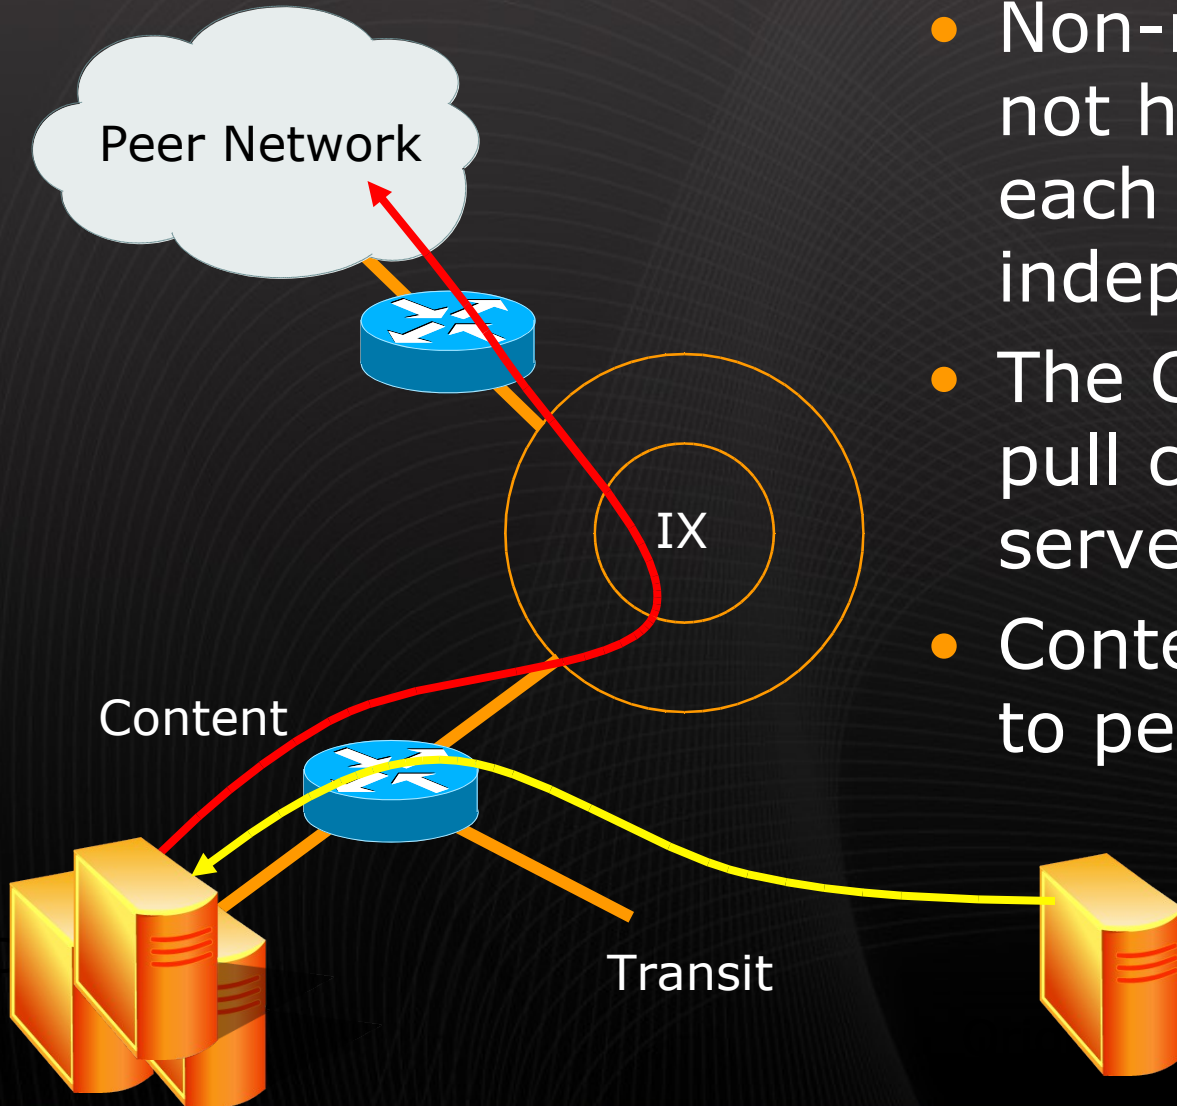
- Performance
 - CDNs and ISPs are in the same business, just on different sides - we both want to serve end users as quickly and reliably as possible
 - You know more about your network than any CDN ever will, so working with the CDN directly can help them deliver the content more quickly and reliably
- Cost Reduction
 - Transit savings
 - Possible backbone savings

Why ISPs peer with CDNs



- Marketing
 - Claim performance benefits over competitors
 - Keep customers from seeing “important” web sites through their second uplink
- Because you are nice :-)

How Non-Network CDNs use IXes



- Non-network CDNs do not have a backbone, so each IX instance is independent
- The CDN uses transit to pull content into the servers
- Content is then served to peers over the IX

How CDNs use IXes



- Non-network CDNs usually do not announce large blocks of address space because no one location has a large number of servers
 - It is not uncommon to see a single /24 from a CDN at an IX
- This does not mean you will not see a lot of traffic
 - How many web servers does it take to fill a gigabit these days?

Questions?



- patrick@akamai.com