

# The forces behind the changing Internet: IXPs and content delivery and SDN

Steve Uhlig

Queen Mary, University of London

[steve@eecs.qmul.ac.uk](mailto:steve@eecs.qmul.ac.uk)

<http://www.eecs.qmul.ac.uk/~steve/>

Credit to collaborators:

Bernhard Ager, Nikos Chatzis, Anja Feldmann, Benjamin Frank, Bruce Maggs, Wolfgang Mühlbauer, Ingmar Poese, Nadi Sarrar, Georgios Smaragdakis, Walter Willinger

# Agenda

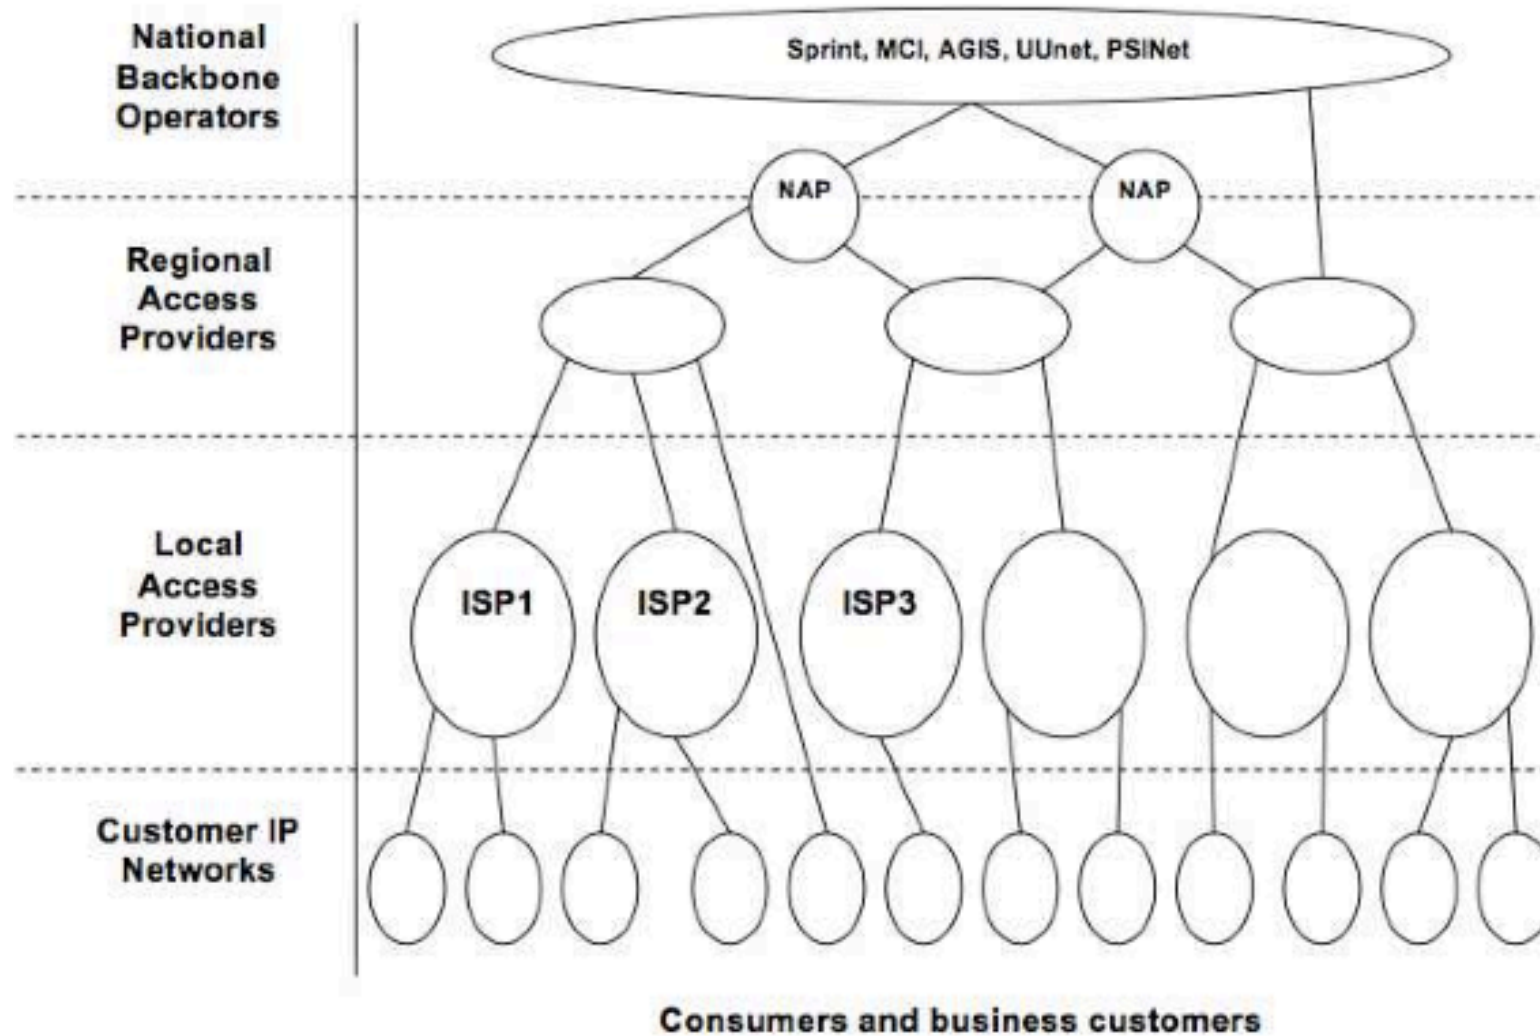
- **Internet update**
- Internet exchange points
- Content delivery ecosystem
- Software Defined Networking

# Accepted view of the Internet

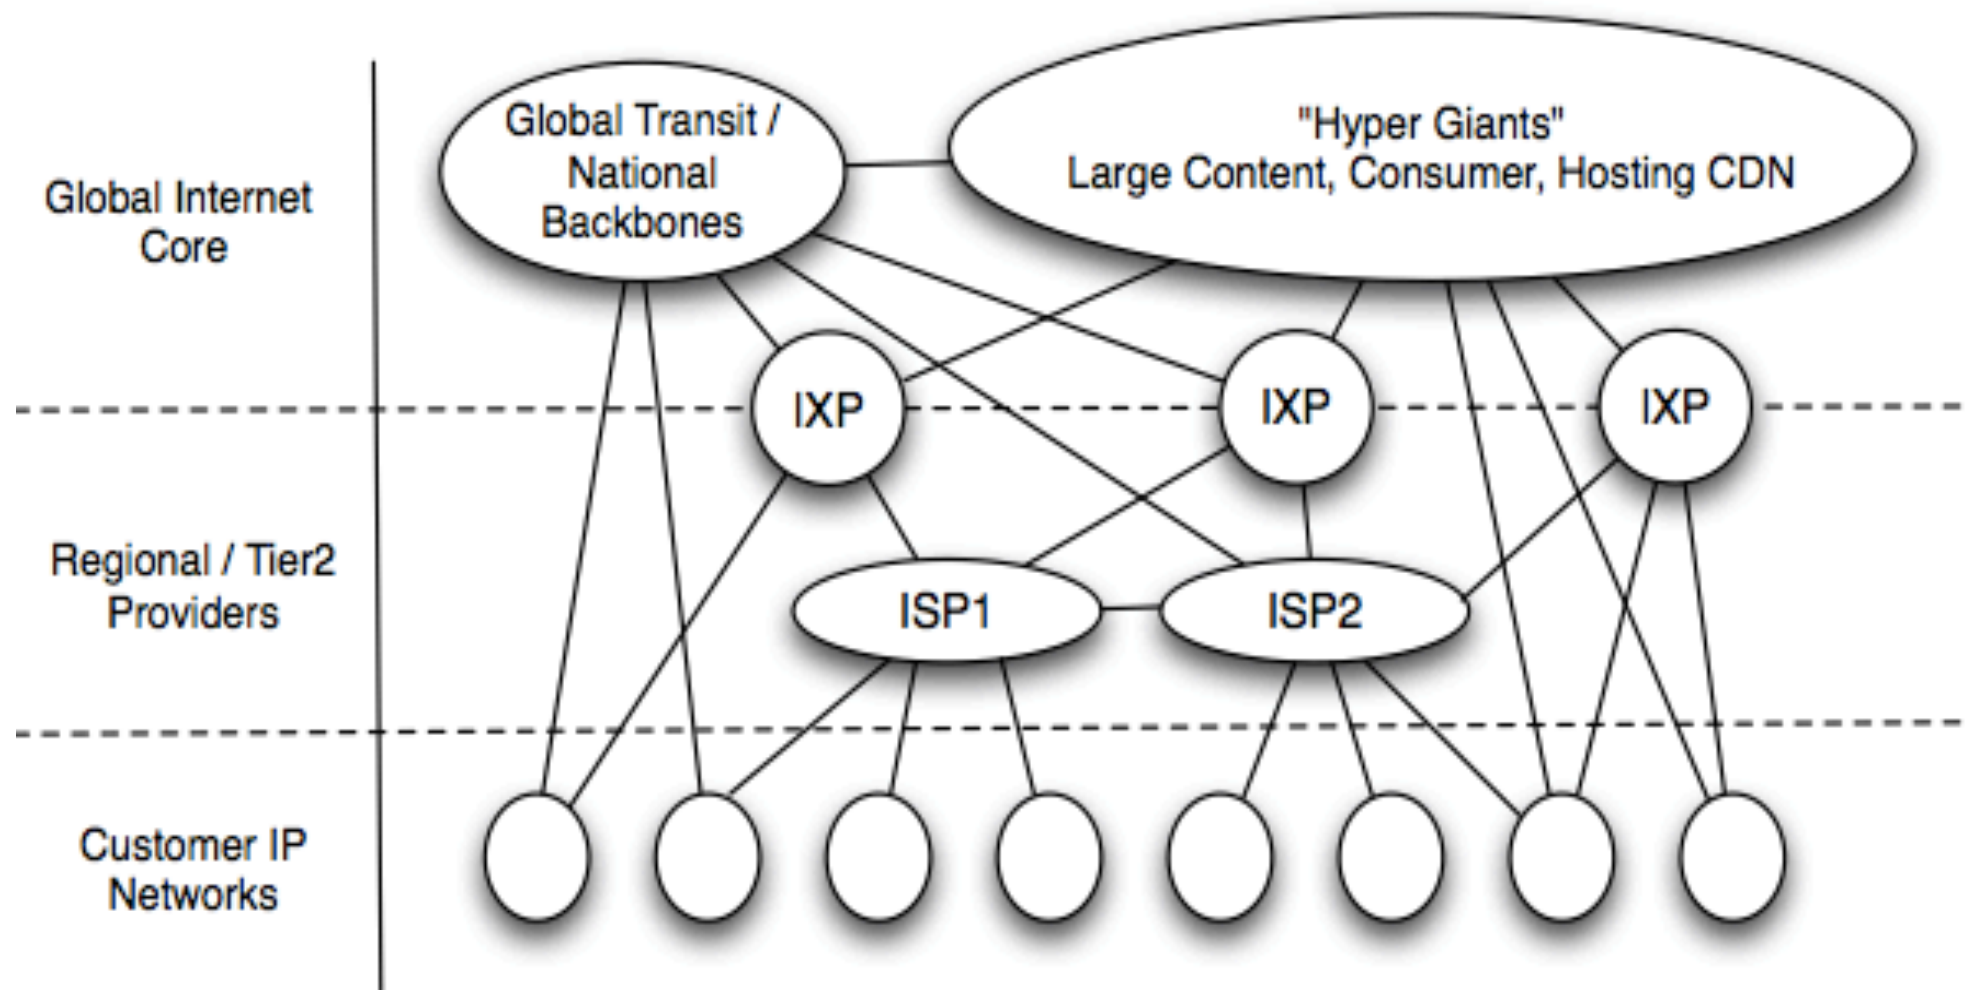


- 35,000+ networks
- Hierarchical structure
  - Tier-1 (10-20): ATT, L3, Sprint,...
  - Regional ISPs (15%): BT, Telefonica,...
  - Stubs (85%): eyeball ISPs, universities, enterprise networks
- Known AS connectivity
  - Customer-provider: 90,000+
  - Peer-peer: 40,000?

# Old mental model



# Most recent mental model



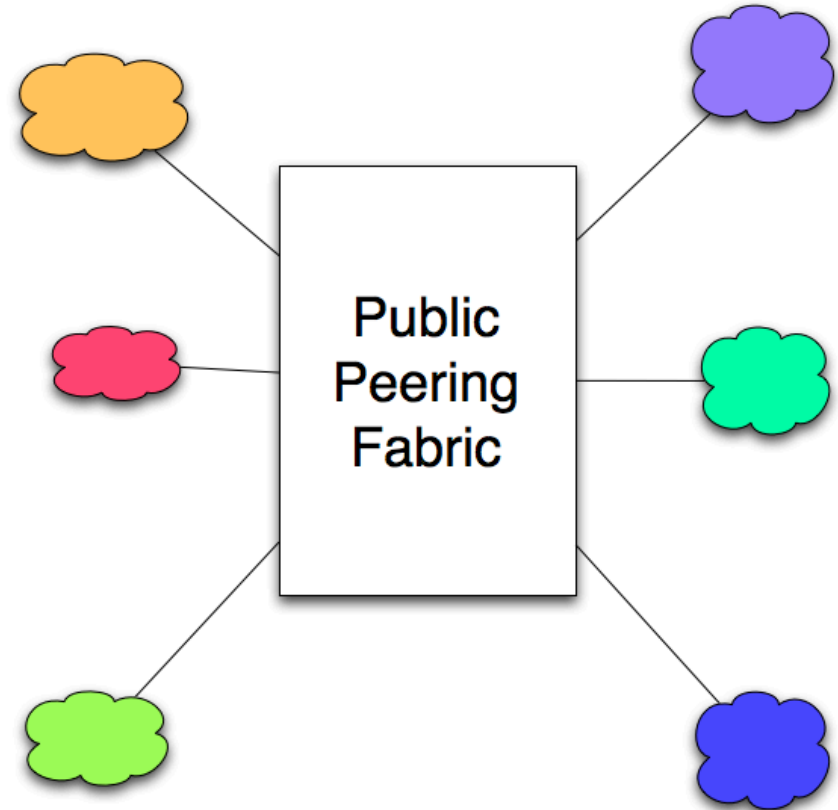
C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. *Internet Interdomain Traffic*. Proc. of ACM SIGCOMM 2010.

# Agenda


- Internet update
- **Internet Exchange Points**
- Content delivery ecosystem
- Software Defined Networking

# Internet Exchange Point

- An **Internet exchange point (IXP)** is a layer 2 service to facilitate the interconnection between
  - ISPs
  - Hosting or service providers
  - CDNs
- An IXP facilitates peering between players, usually across a public and/or private peering fabric of some type
- Offer public and/or private peerings



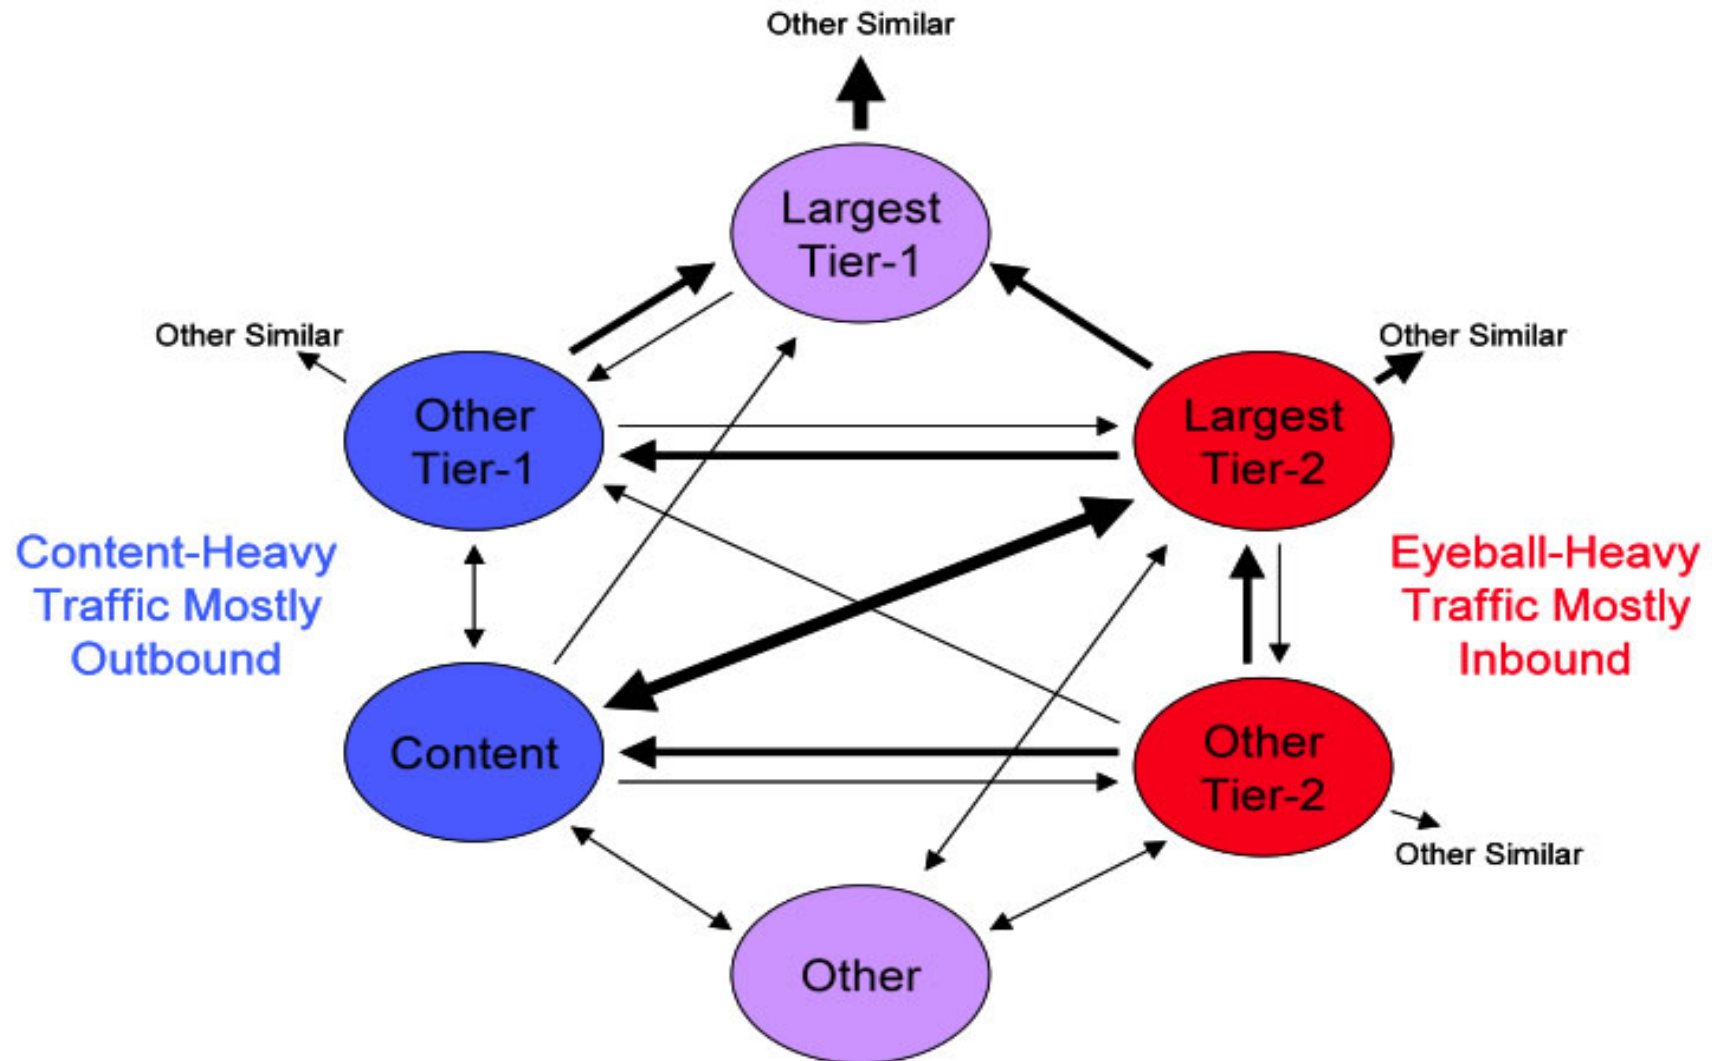
# IXP architecture: example



B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. *Anatomy of a Large European IXP*. Proc. of ACM SIGCOMM 2012.




# Interconnection and business



William B. Norton. *The Internet Peering Playbook : Connecting to the Core of the Internet*. DrPeering Press, 2012.

# Peerings at an IXP



- Mixed structure
  - Tier-1 (10-20) + Large IXPs
  - Regional ISPs + smaller regional IXP
- Known AS connectivity
  - Customer-provider: 90,000+
  - Peer-peer: data from a single IXP doubles it!
- Traffic
  - Increasingly exchanged directly between CDN and regional ISPs

# Agenda

- Internet update
- Internet Exchange Points
- **Content delivery ecosystem**
- Software Defined Networking

# World data centers



<http://www.datacentermap.com/>

# Google data centers



<http://royal.pingdom.com/2008/04/11/map-of-all-google-data-center-locations/>

# World clouds



<http://www.datacentermap.com/>

# Today's popular CDNs

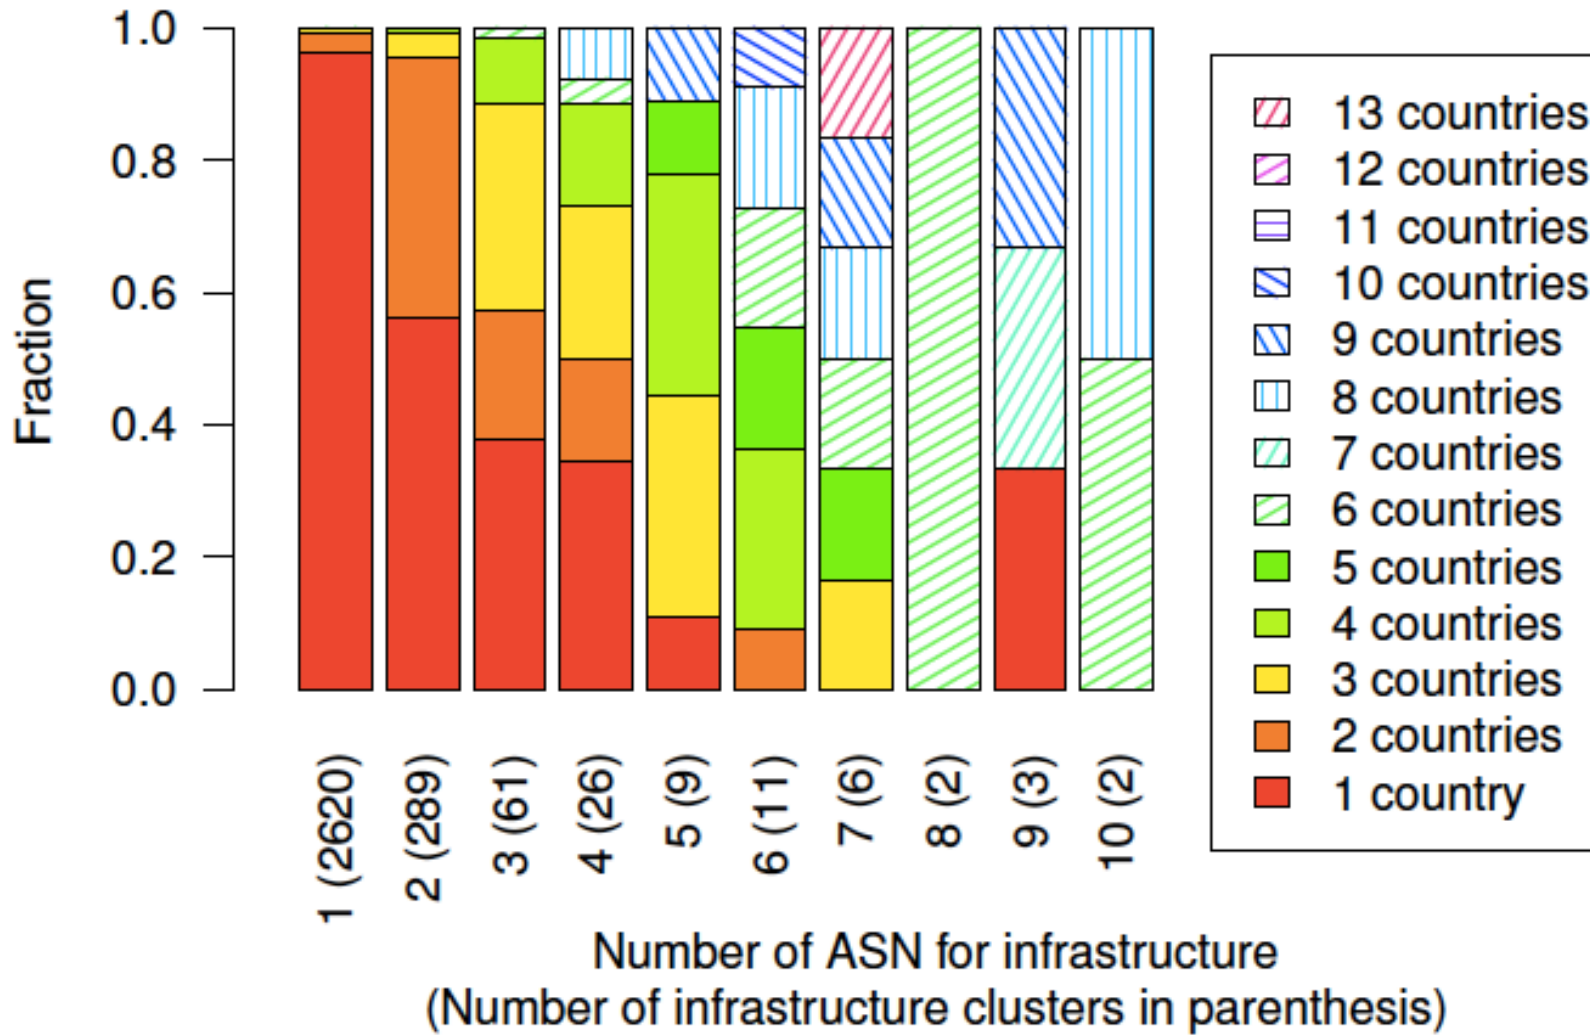
Rank	# hostnames	# ASes	# prefixes	owner	content mix
1	476	79	294	Akamai	
2	161	70	216	Akamai	
3	108	1	45	Google	
4	70	35	137	Akamai	
5	70	1	45	Google	
6	57	6	15	Limelight	
7	57	1	1	ThePlanet	
8	53	1	1	ThePlanet	
9	49	34	123	Akamai	
10	34	1	2	Skyrock	
11	29	6	17	Cotendo	
12	28	4	5	Wordpress	
13	27	6	21	Footprint	
14	26	1	1	Ravand	
15	23	1	1	Xanga	
16	22	1	4	Edgecast	
17	22	1	1	ThePlanet	
18	21	1	1	ivwbox.de	
19	21	1	5	AOL	
20	20	1	1	Leaseweb	

■ only on TOP, ■ both on TOP and EMBEDDED, ■ only on EMBEDDED, ■ TAIL.

B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig. *Web content cartography*.  
Proc. of ACM SIGCOMM IMC 2011.



# CDNs geographic coverage



# Content is power


Rank	CAIDA-degree	CAIDA-cone	Renesisys	Knodes	Arbor	Potential	Normalized potential
1	Level 3	Level 3	Level 3	Level 3	Level 3	NTT	Chinanet
2	Cogent/PSI	AT&T	Global Crossing	Cogent	Global Crossing	Tinet	Google
3	AT&T	MCI	Sprint	Global Crossing	Google	Global Crossing	ThePlanet
4	MCI	Cogent/PSI	NTT	Sprint	*	Deutsche Telekom	SoftLayer
5	Hurricane	Global Crossing	Savvis	Tinet	*	KDDI	China169 backbone
6	Qwest	Sprint	TeliaSonera	NTT	Comcast	Telia	Level 3
7	Sprint	Qwest	Tinet	AT&T	*	Akamai	Rackspace
8	Global Crossing	Hurricane Electric	Verizon	Swisscom	*	Bandcon	China Telecom
9	tw telecom	tw telecom	AT&T	Hurricane	*	Cable and Wireless	1&1 Internet
10	INIT7	TeliaNet	China Telecom	Telia	*	Qwest	OVH

- CAIDA: BGP-based degree or customer-cone
- Renesisys: variant of CAIDA-cone
- Knodes: Fixedorbit.com centrality metric
- Arbor: Interdomain traffic
- Potential: hostname-based
- Normalized potential: weighted hostnames

B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig. *Web content cartography*. Proc. of ACM SIGCOMM IMC 2011.

# CDN 3.0

- Hybrid infrastructures: Akamai, PPTV
- Meta-CDNs, e.g., Conviva
- Virtual CDNs through ISP micro-datacenters



# Agenda

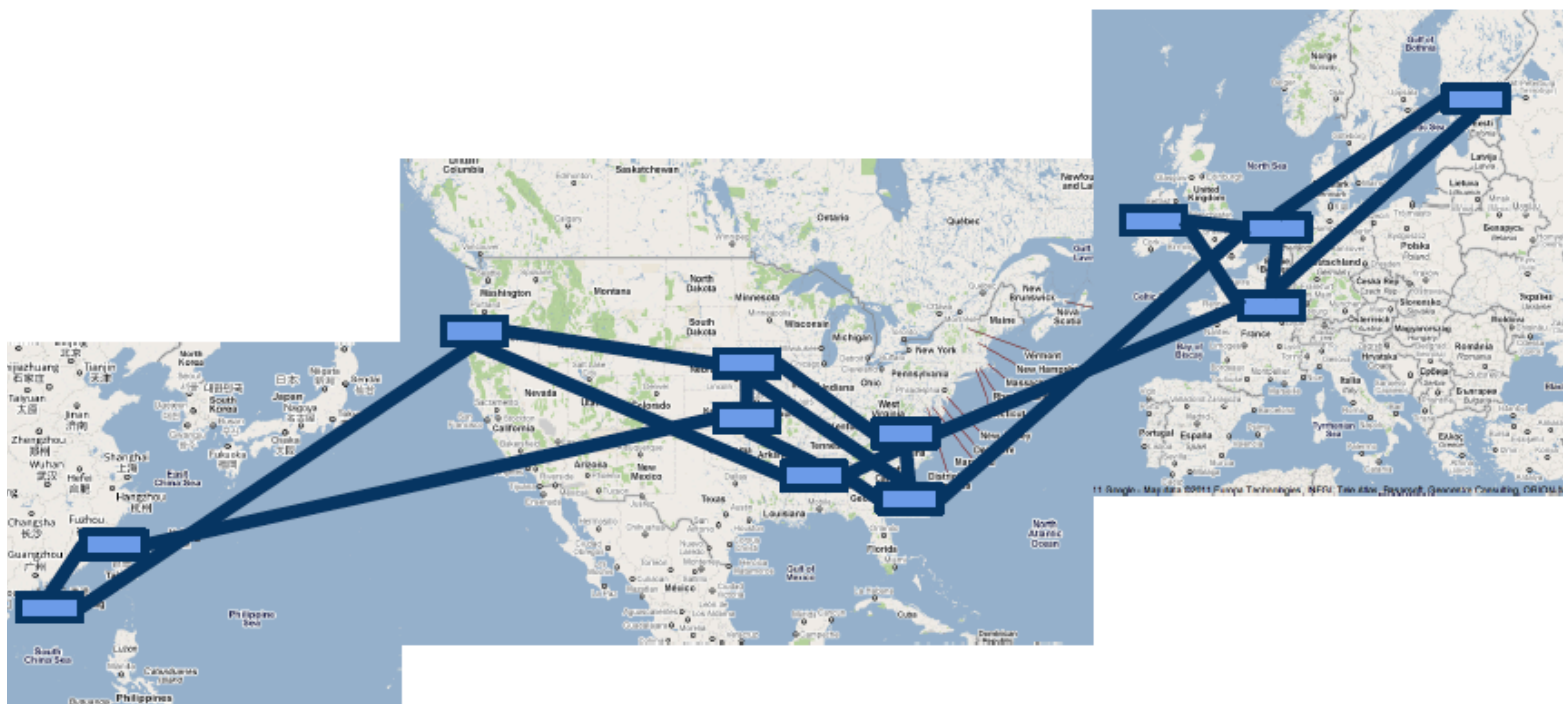
---

- Internet update
- Internet exchange points
- Content delivery ecosystem
- **Software Defined Networking**

# Google and SDN

- Google is using OpenFlow
- Purpose: traffic engineering

## Google's Software Defined WAN



# Conclusion

- Revisiting mental model of the Internet
- IXP has a rich ecosystem: members, traffic, geography
- CDN 3.0: hybrid and ISP-CDN clouds
- SDN: TE as a use-case