Lecture 16: Internet Peering

CS 234 / NetSys 210: Advanced Computer Networks

Sangeetha Abdu Jyothi

UCIrvine

This lecture uses material from IMC'13 talk on IXP, slides from Anja Feldman, Phillippa Gill and Morley Mao.

BGP route propagation

- Connectivity does not imply reachability
- Not all possible routes propagate
- Commercial relationships determine policies for
    - Route import
    - Route selection
    - Route export
- Typical relationships
    - Provider-customer: customer pay money for transit
    - Peer-peer: typically exchange respective customers' traffic for free

AS Relationships

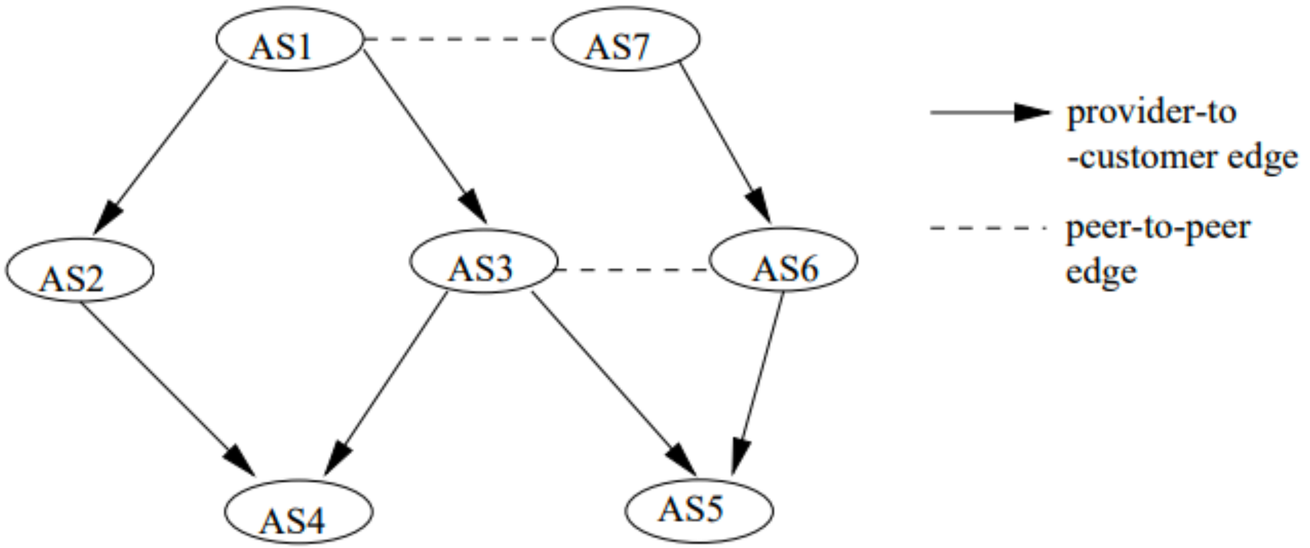

Effects of Changing AS relationships

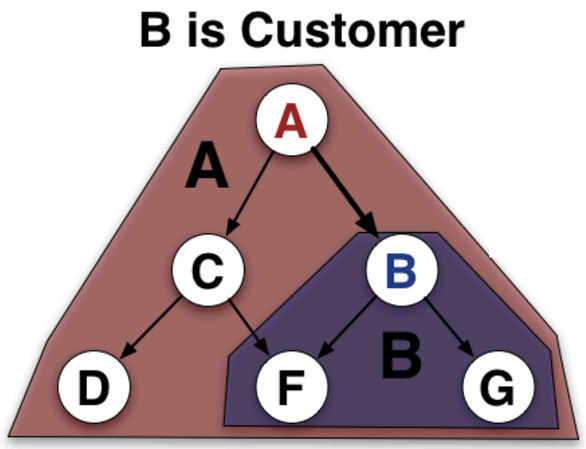

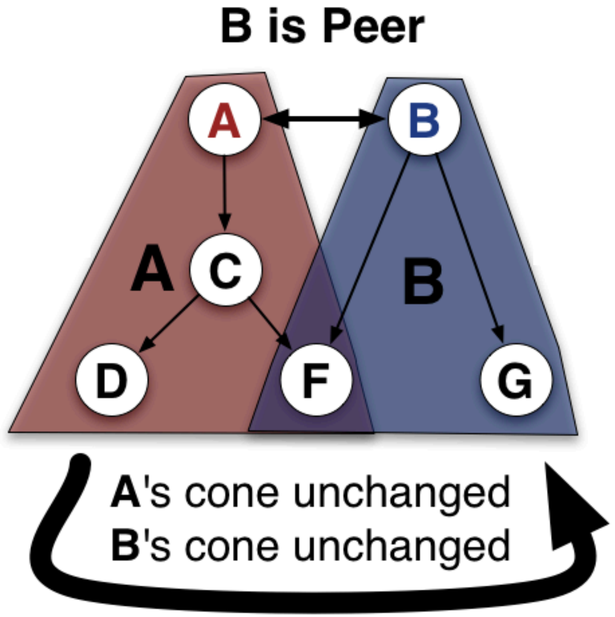

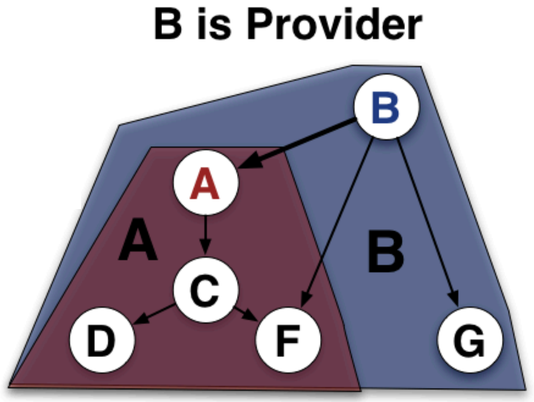

How does Peering Relationships help?

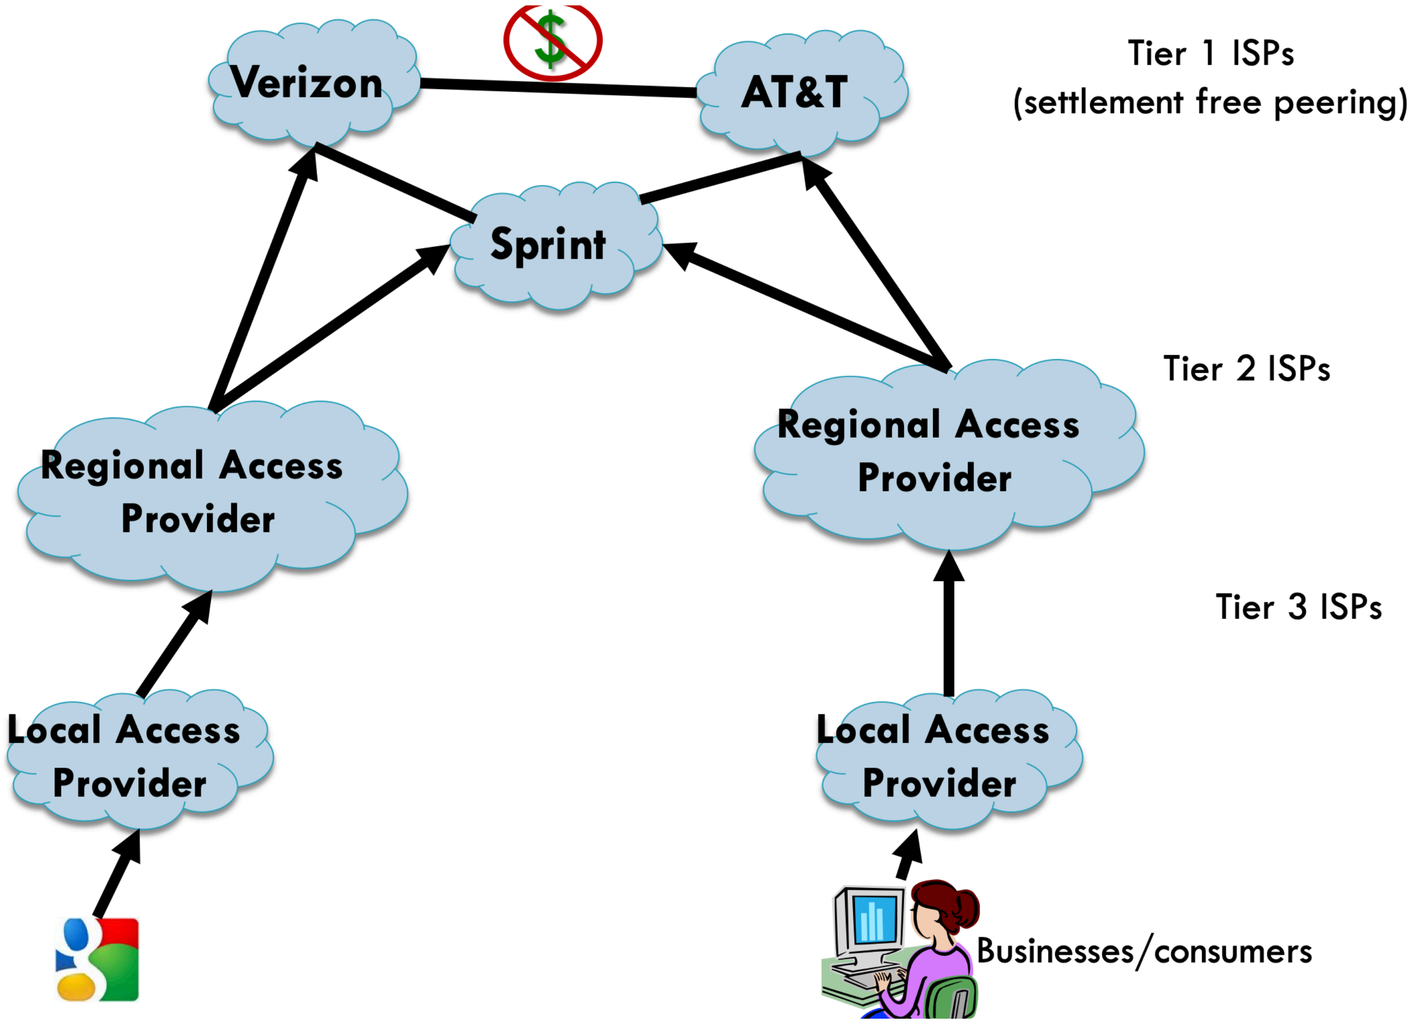

How does Peering Relationships help?

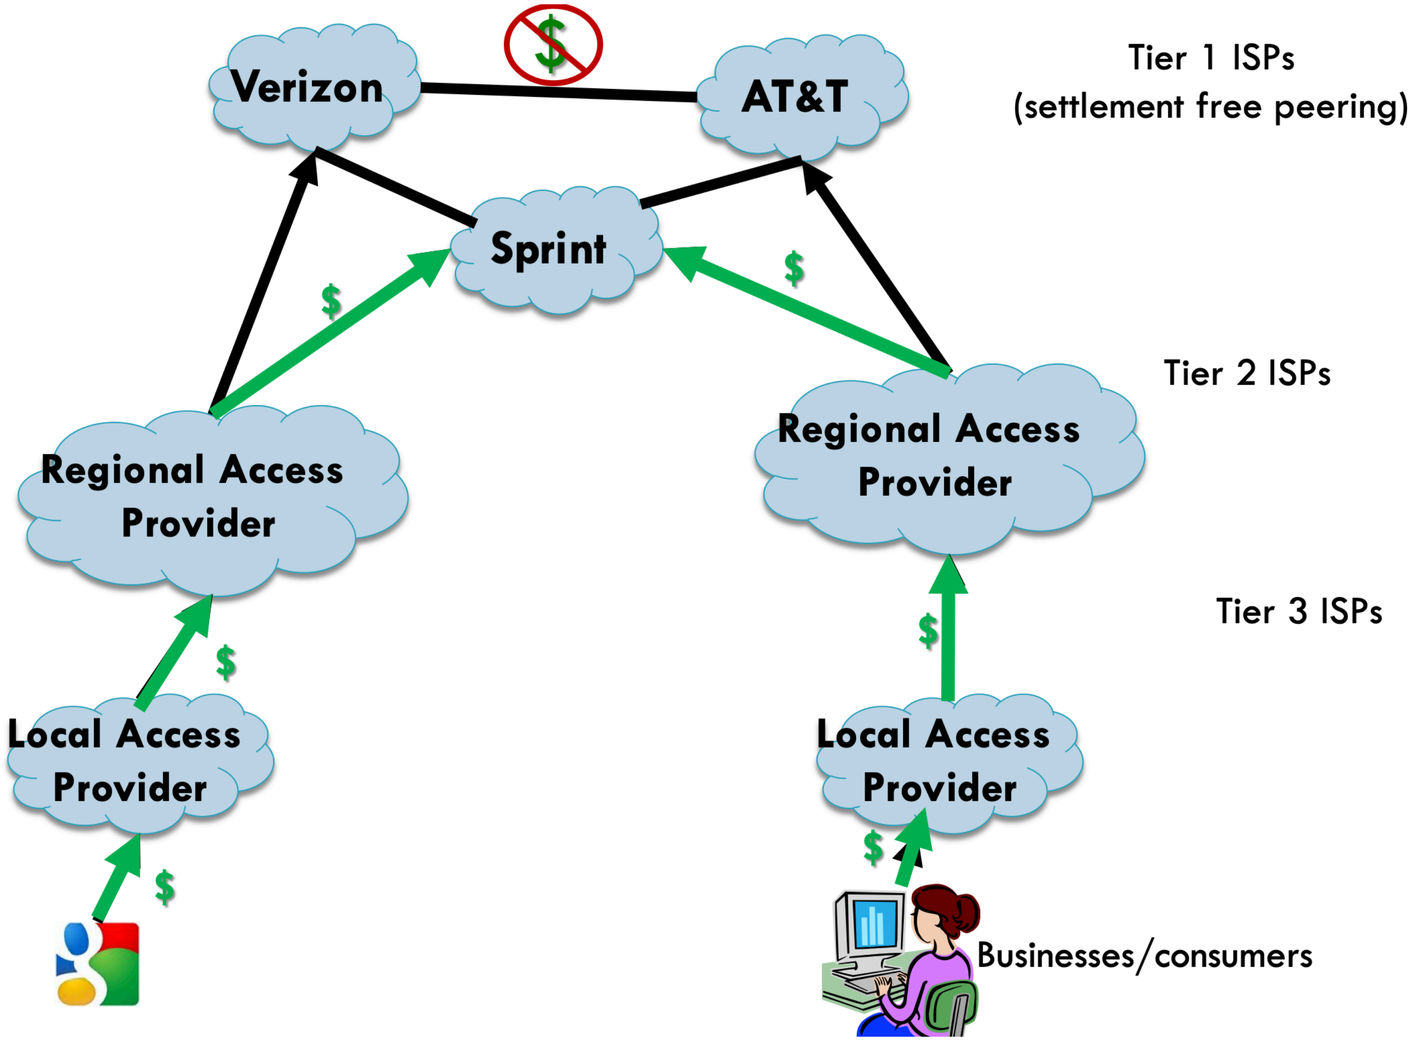

How does Peering Relationships help?

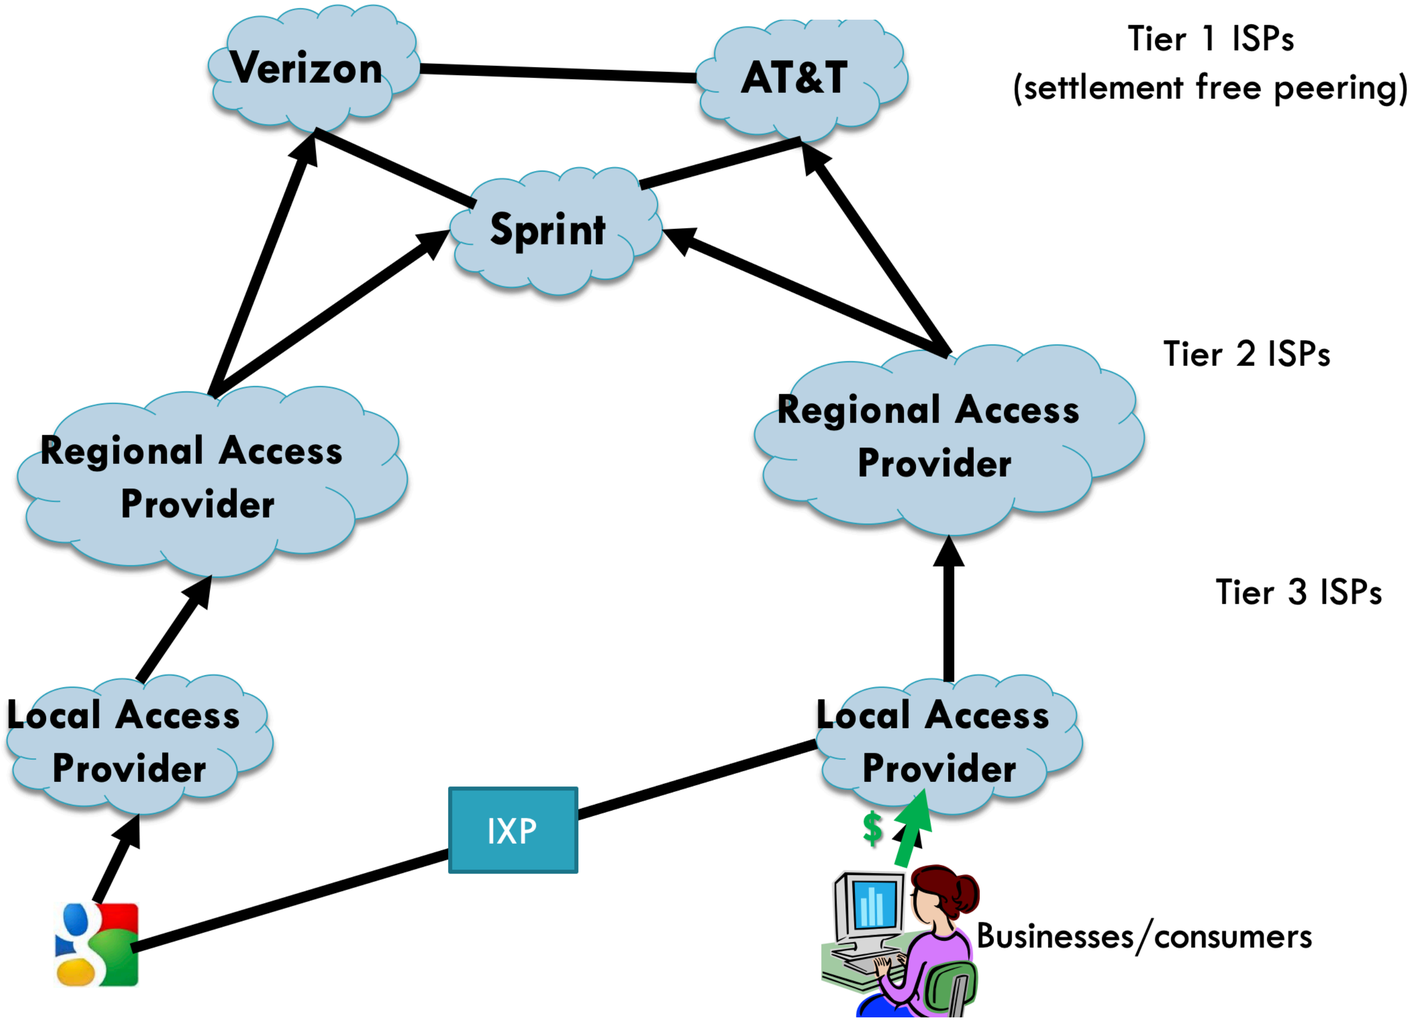

How does Peering Relationships help?

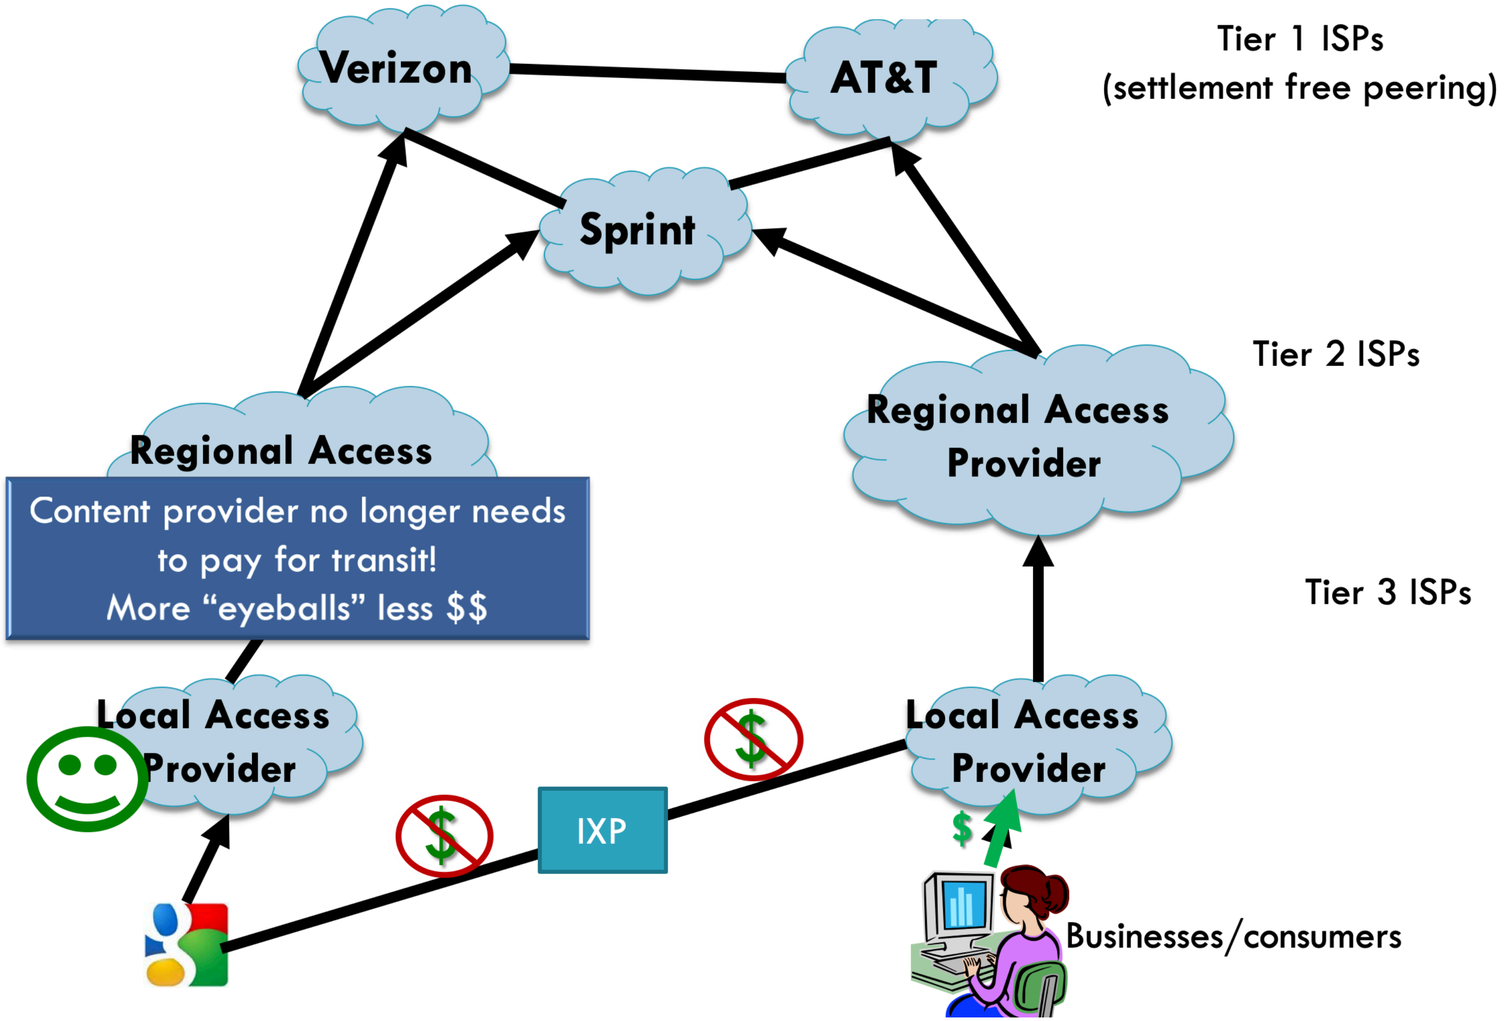

How does Peering Relationships help?

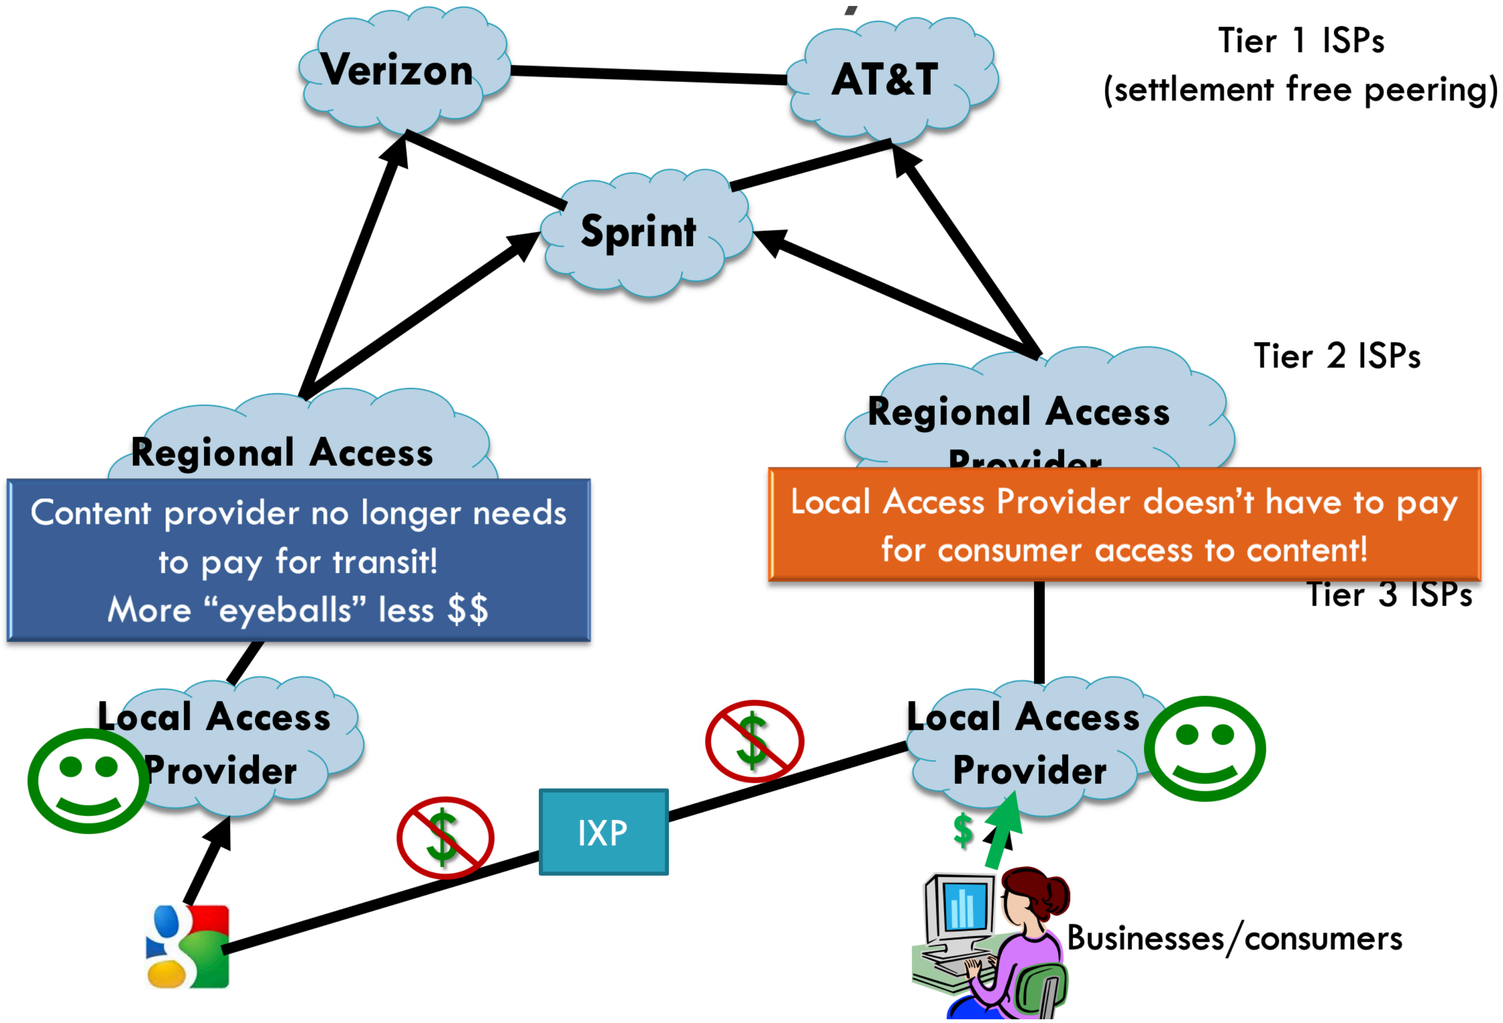

Traditional Internet Model

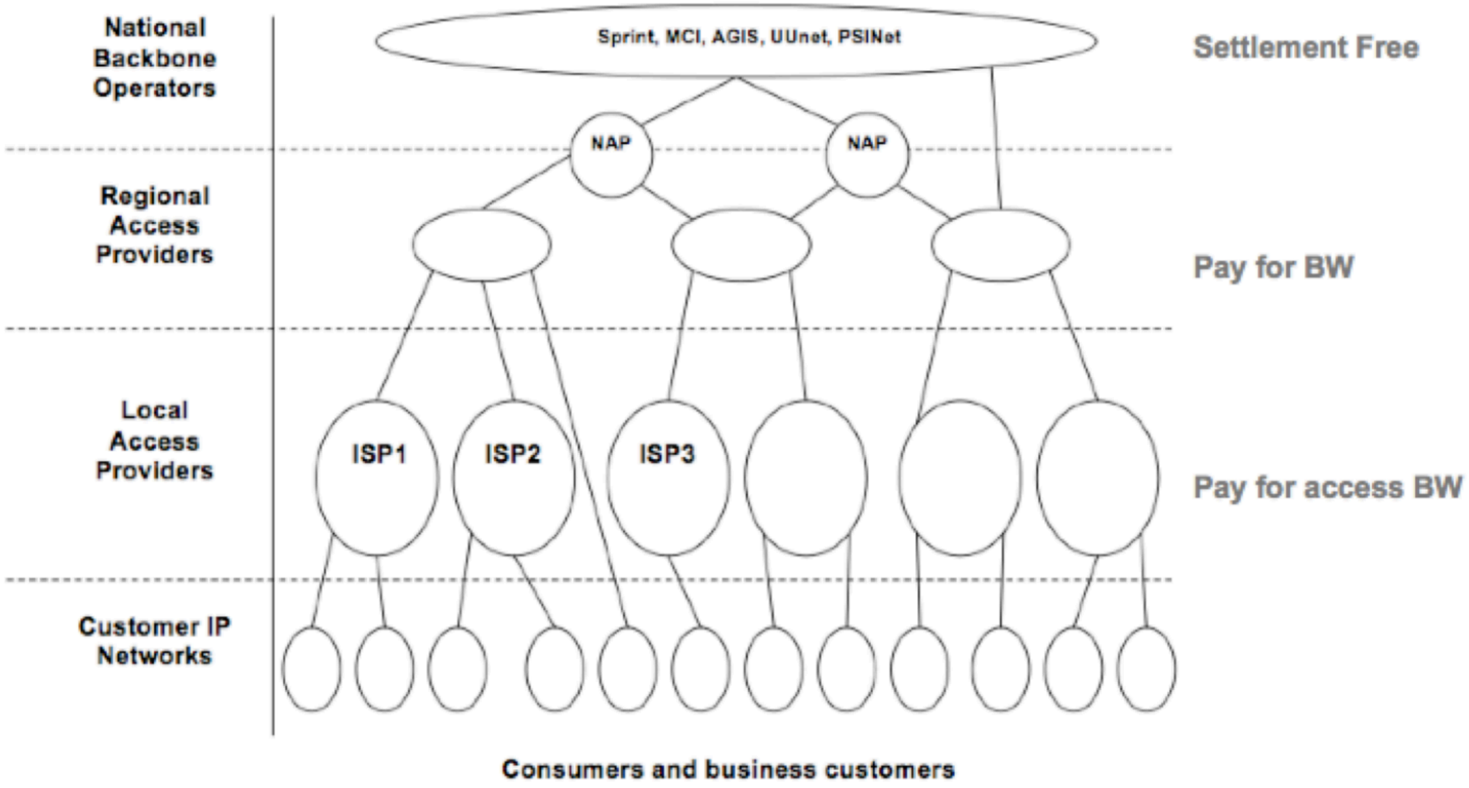

The new Internet model

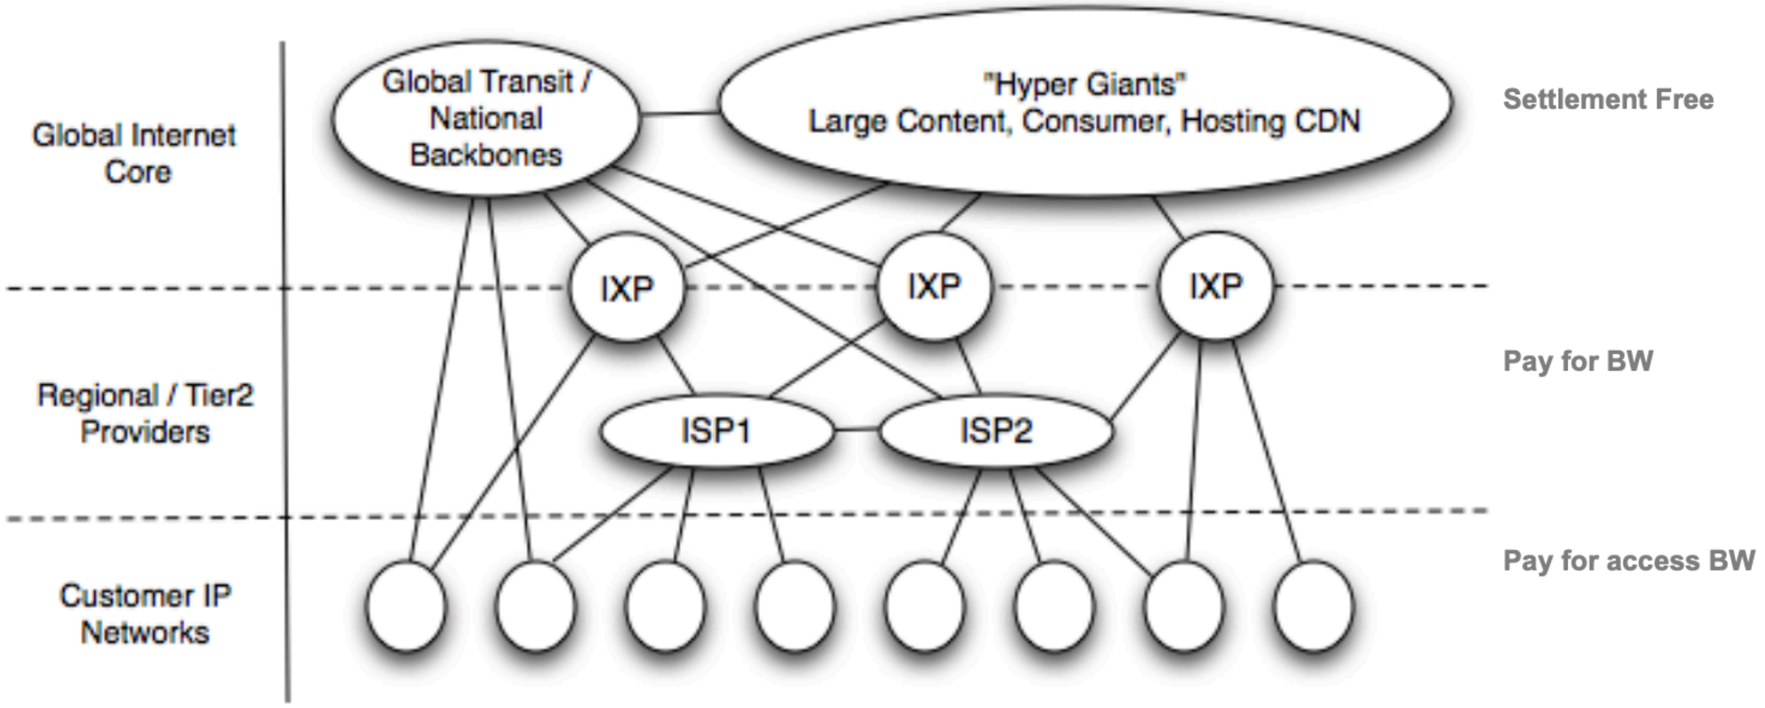

Tier-1 vs Tier-2 peering

- Tier 1 ISPs
    - Buy no transit from any other providers
    - Have only customers and peers
    - Has full mesh peering with other Tier 1's
- Tier 2 ISPs
    - ISP that purchases (resells) transit within an Internet region

Benefit of tier-2 peering

- Decreases the cost and reliance on purchased Internet transit
- Lowers inter-AS traffic latency
    - Fewer AS hops, AS peering links traversed
- Is peering always better than transit?
- Concerns of peering:
    - Traffic asymmetry
    - No SLAs: less liability or incentive to improve performance
    - “free” rather than getting paid

Peering Wars

- Reasons to Peer
  - Reduces upstream transit costs
  - Can increase end-to-end performance
- Reasons to Not Peer
  - You would rather have customers
  - Peers are usually your competition
  - Peering relationships may require periodic renegotiation

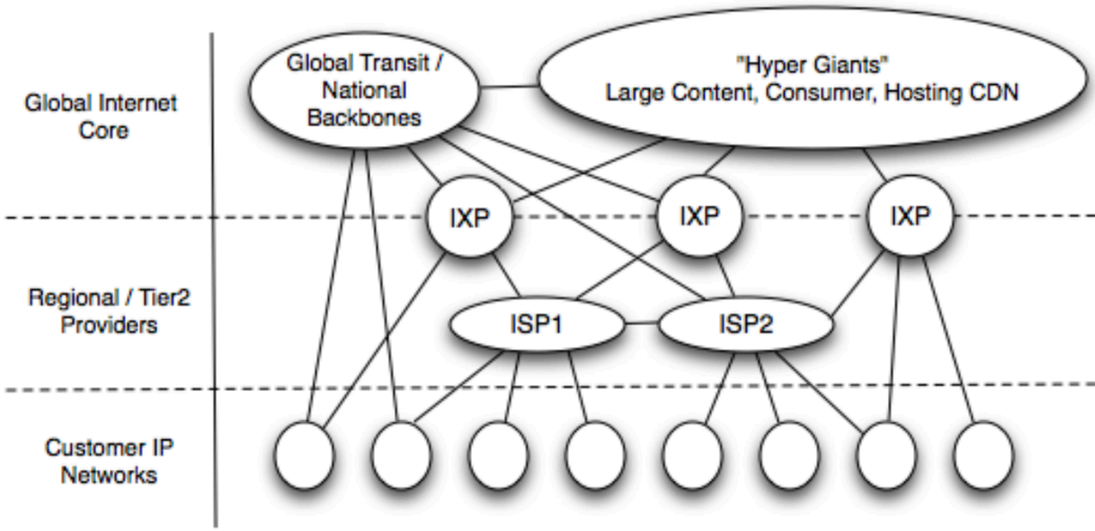

IXPs

Physical locations that offer a shared (often distributed) layer-2 switching fabric for members (networks) to exchange traffic with one another.

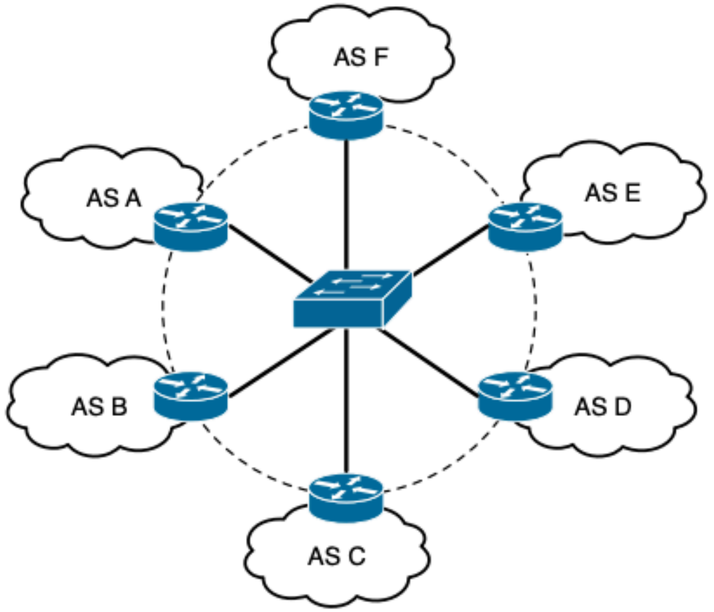

IXP in Reality is more than a switch

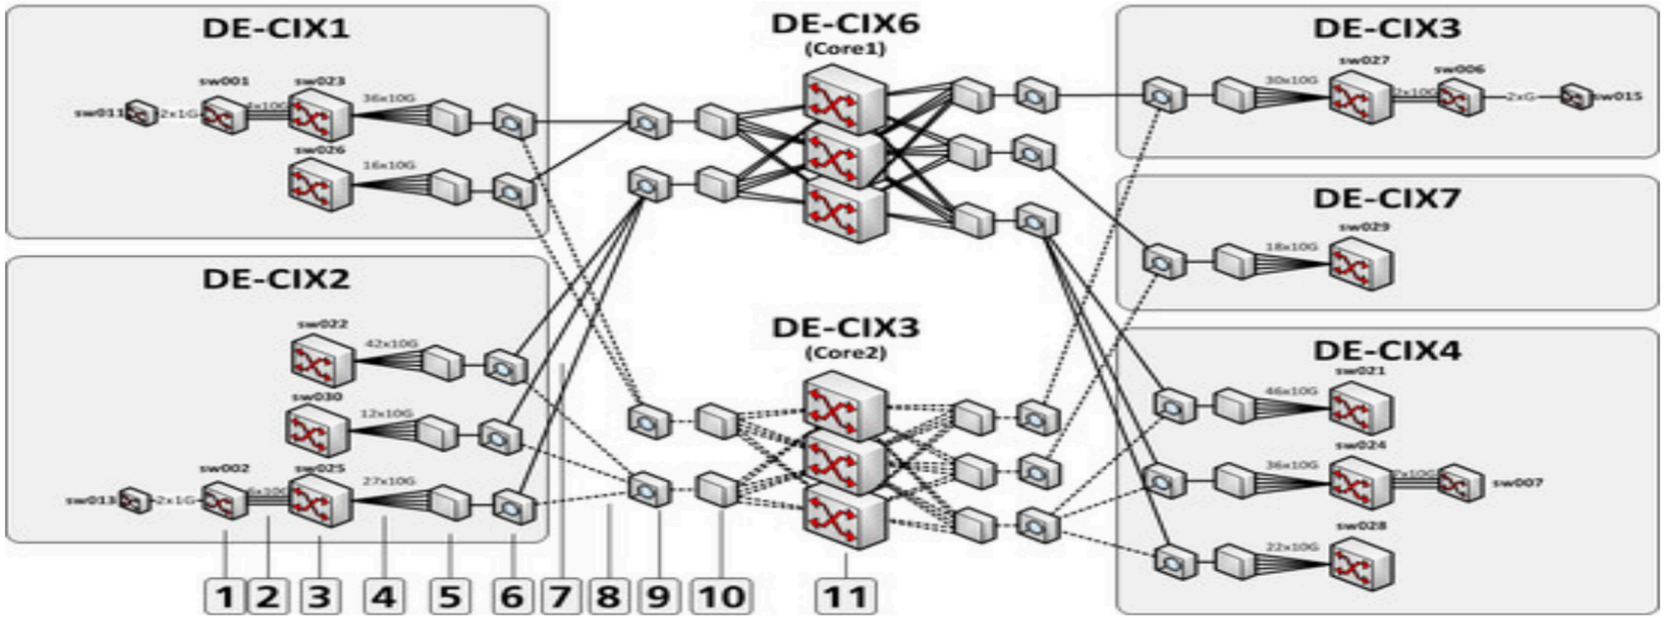

- Complex system
- A number of services are offered

Global Map of IXPs

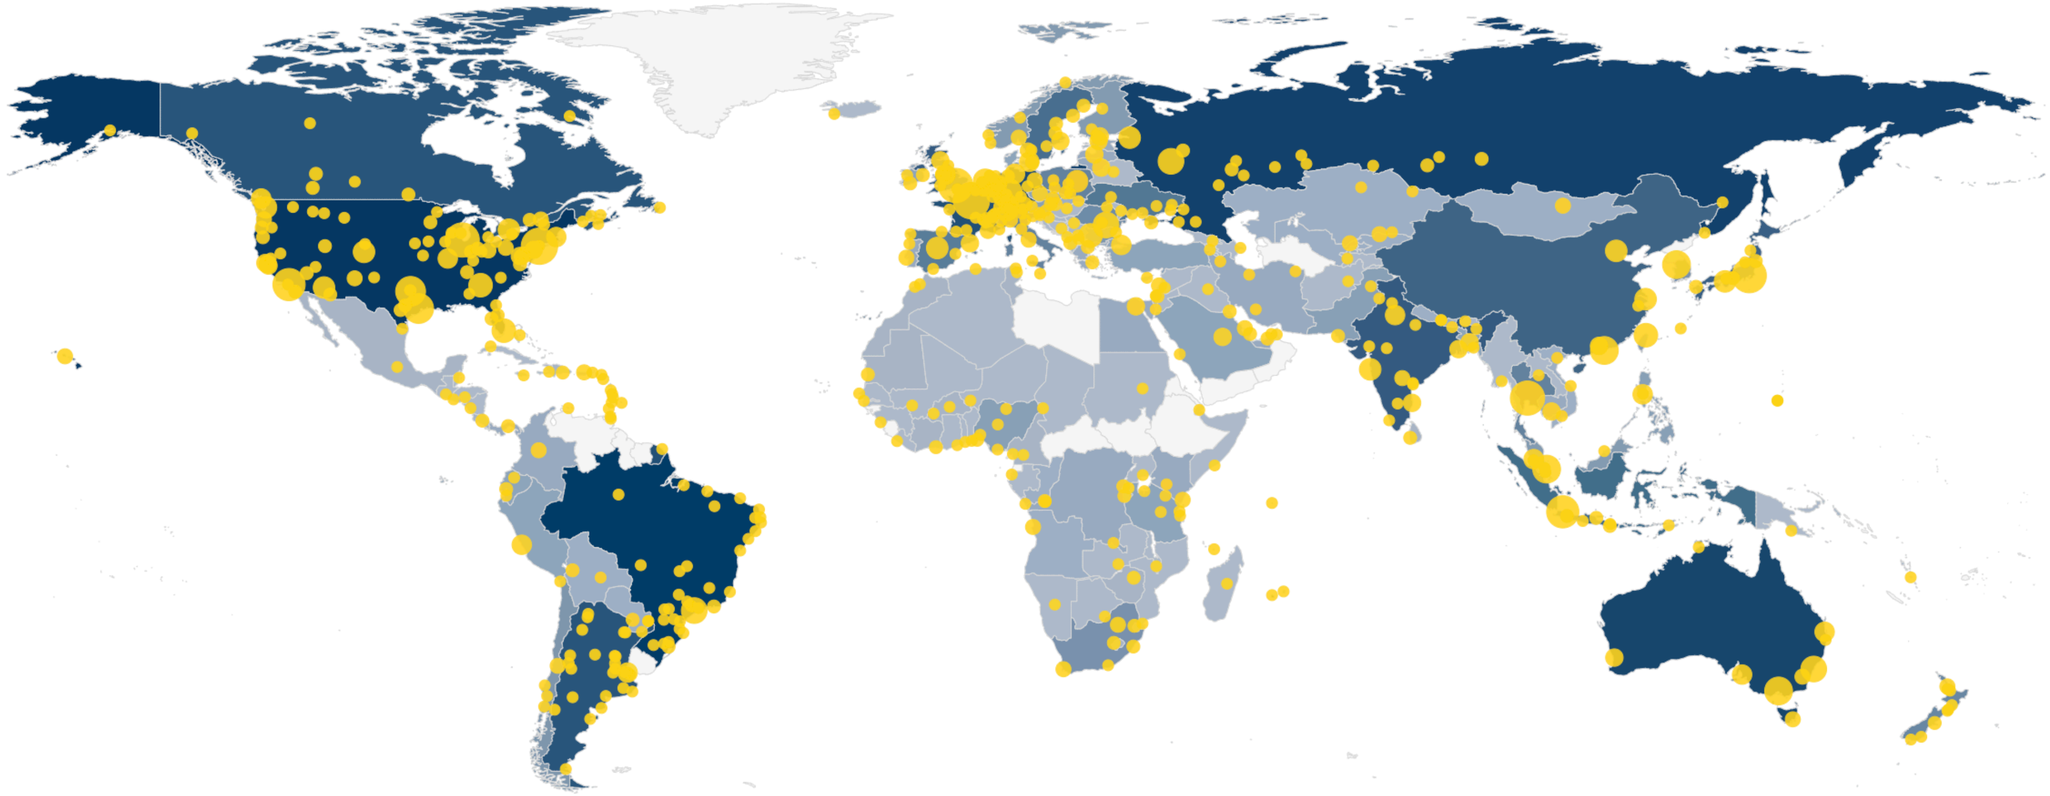

Image from Packet Clearing House

Anatomy of a Large European IXP

- sFlow data
  - Sampling Rate: 1/16K packets
  - Sampling Size: First 128 bytes of Ethernet frame

Table 1: Overview of IXPs sFlow dataset.

	Apr 25 May 1	Aug 22 Aug 28	Oct 10 Oct 16	Nov 28 Dec 4
Identified member ASes	358	375	383	396
Router IPs	426	445	455	474
MAC addresses	428	448	458	474
Tier-1	13	13	13	13
Tier-2	281	292	297	306
Leaf	64	70	73	77
Countries of member ASes	43	44	45	47
Continents of member ASes	3	3	3	3
Average packet rate (Mpps)	142	150	166	174
Average bandwidth (Gbps)	838	863	954	992
Daily avg volume (PB)	9.0	9.3	10.3	10.7

A single IXP had more P2P links than inferred datasets

- 50,000 P2P links in the dataset!

Table 2: Overview of routing and looking glass datasets for November. The numbers show P-P links.

Dataset	Unique LGs / ASN	Visible links	only in this dataset
RV	78	5,336	1,084
RIPE	319	10,913	5,460
NP	723	3,419	684
RV+RIPE+NP	997	13,051	10,472
LG	821 / 148	4,892	2,313
RV+RIPE+NP+LG	1,070	15,364	15,364

Number of Peers

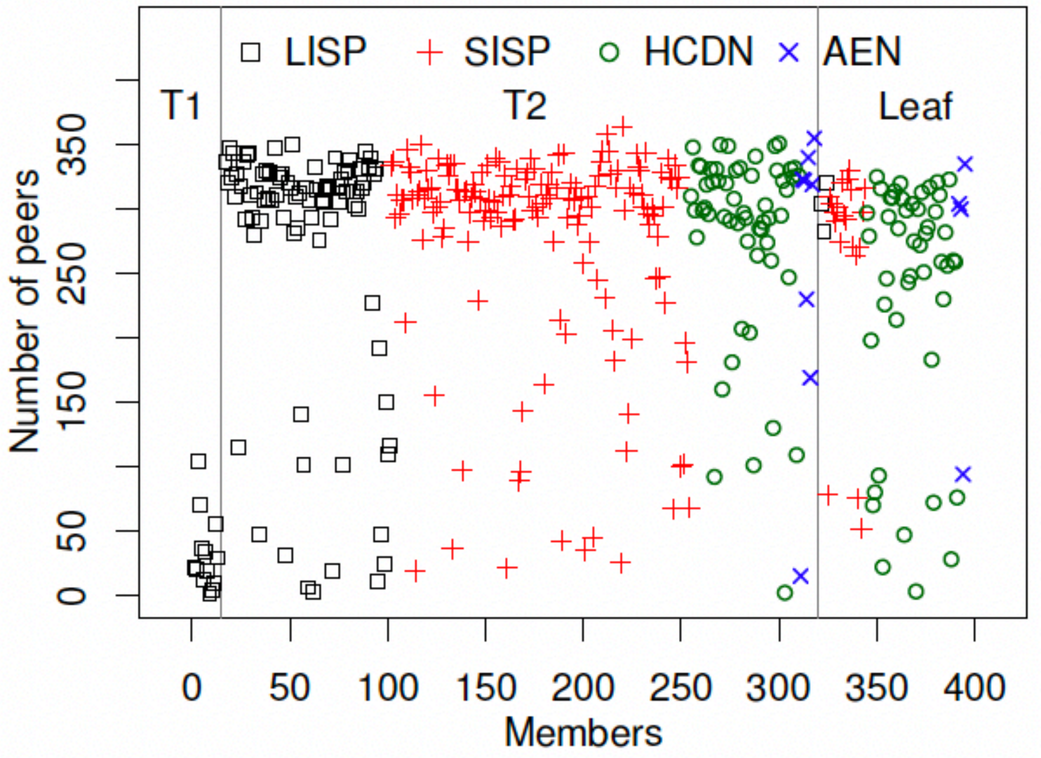

(b) Scatter-plot of num. of peers per member.

Traffic Asymmetry

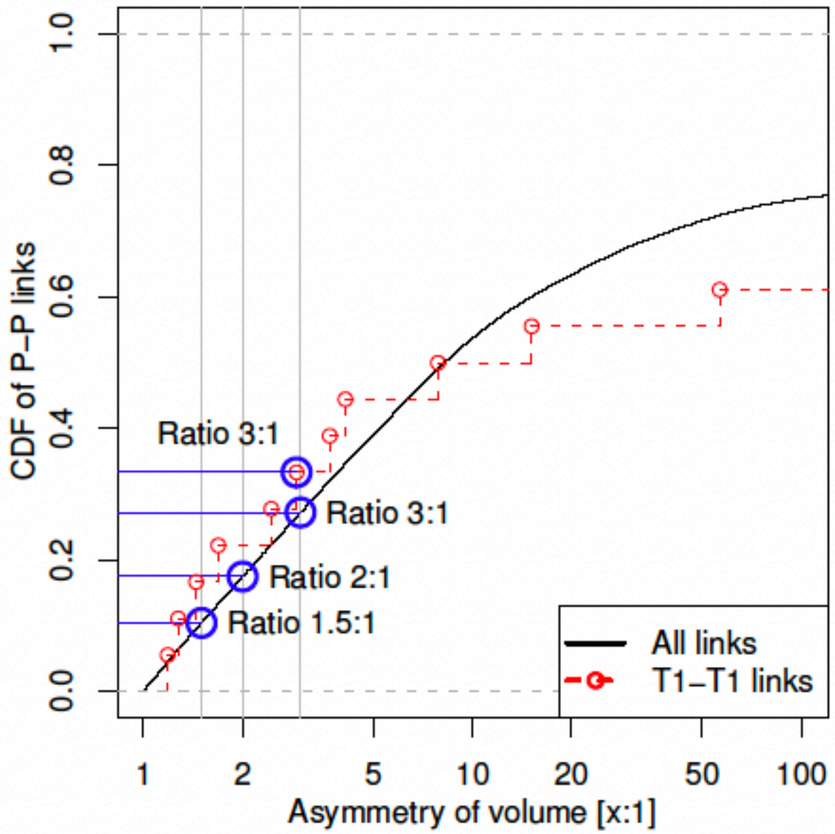

(a) Traffic asymmetry across P-P links.

Traffic Destinations

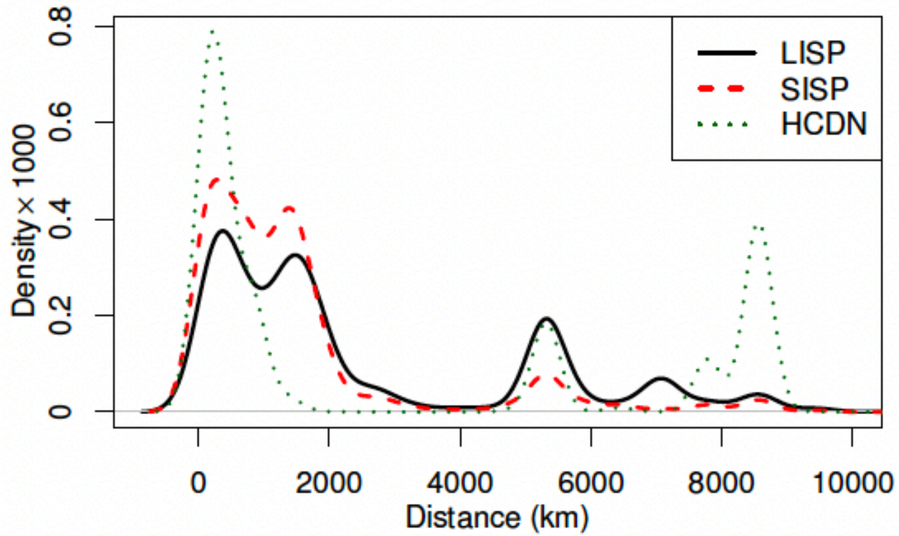

Figure 6: Geographic distances of IP endpoints to IXP.

Recent Statistics - DE-CIX Frankfurt

Traffic Frankfurt – 2 days

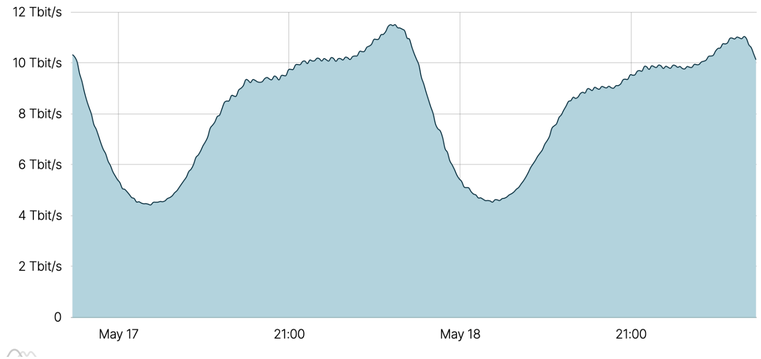

ALL-TIME PEAK
11.97 Tbit/s

GRAPH PEAK
11.53 Tbit/s

GRAPH AVERAGE
9.05 Tbit/s

CURRENT
10.15 Tbit/s

ACTIVE ASNs	CONNECTING ASNs	IPv4 ROUTES	IPv6 ROUTES
1071	9	650,445	97,610

Traffic Frankfurt – 5 years

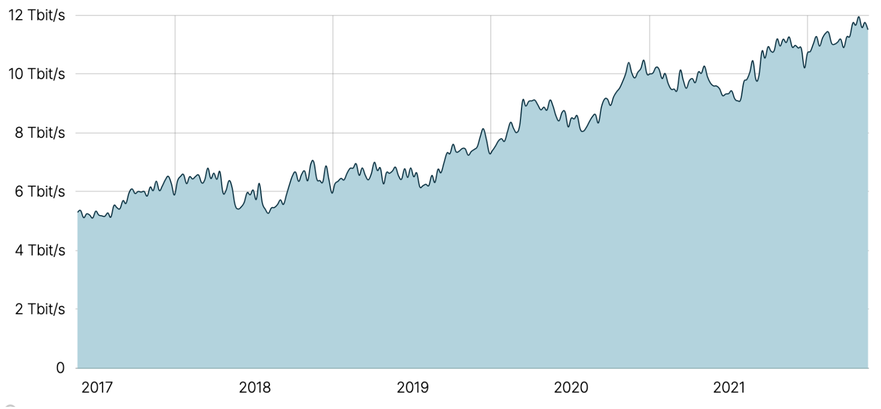

ALL-TIME PEAK
11.97 Tbit/s

GRAPH PEAK
11.97 Tbit/s

GRAPH AVERAGE
5.29 Tbit/s

CURRENT
10.15 Tbit/s

Techniques for Internet Topology Discovery

- Interface Level
    - Traceroute
    - Geolocation
        - Delay based techniques
        - Name to location mapping
- Router Level
    - Subnet discovery
- AS level
    - BGP routes
    - Traceroute + IP to AS mapping

Thanks!