

Anatomy of a large european IXP


Bernhard Ager • Nikolaos Chatzis • Anja Feldmann
Nadi Sarrar • Steve Uhlig • Walter Willinger
Proc. of ACM SIGCOMM '12, 42(4):163-174, 2012.

Presented by Han Zhang and Ryan Marcotte


What is an IXP?

Internet Exchange Point

- Physical infrastructure connecting ASes (ISPs, CDNs, etc.)
- >300 worldwide
- Governed primarily by peering agreements
- Traffic exchange facilitated by Border Gateway Protocol (BGP)



Courtesy Wikimedia Foundation



Why study IXPs?

- Lead to better understanding of:
 - the AS-level Internet
 - peering and its economic considerations
 - Internet inter-domain traffic
- Aid in traffic engineering, protocol design
- Inform decisions for IXPs (e.g new services, infrastructure upgrades, etc.)

This paper cited by 162 others since 2012!

Contributions of the paper

- 1) Rich peering fabric of IXPs
- 2) Traffic analysis: IXPs are a “microcosm of the Internet”
- 3) IXP traffic matrix

Peering in the IXP


Objective: study the peer-to-peer relationships existing between member ASes

- Previous approach: BGP routing information, traceroute measurements
- This approach - sFlow data provided by IXP
 - Peering matrix - who is actually peering with whom?
 - ~78,000 possible P-P links, >50,000 actually established
 - Previous estimates were 35,000-45,000... for the entire Internet!

Peering in the IXP (cont.)


- Visible P-P link
 - Observed in sFlow records and BGP/traceroute data
- Invisible P-P link
 - Observed in sFlow records but not in BGP/traceroute data
- Cannot-Tell P-P link
 - Visible in BGP data, but no traffic exchanged in sFlow records

Peering in the IXP (cont.)



IXP-external datasets miss the vast majority of observed links!


Peering in the IXP (cont.)



Most traffic unaccounted for in IXP-external data!

Diversity of the IXP ecosystem - Peering


- Manually categorize business types for ~400 ASes.
- Tier-1 ISPs are reluctant about peering with (potential) transit customers.
- Large number of peers shows the ease of peering at IXP.




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

Diversity of the IXP ecosystem - Traffic

- Highly skewed - 30% ASes contributes 90% traffic.
- Most ASes has symmetric incoming/outgoing traffic (1:3 to 3:1), while the asymmetry matches different business types, e.g. CDN and ISP.
- Application level traffic follows different distributions depends on business types as well.



(b) Traffic asymmetry (out/in) per member.



Diversity of the IXP ecosystem - Misc

- Most ASes use 10x more IP prefix than the number they server
- Over 70% traffic originates or terminates within 2,000km (1242 mi) radius from the IXP.
 - Slightly larger than the Contiguous US.
- Mathematically proves that we need way less than 12 features to characterize member ASes.
 - Proposes novel and brief way of characterization as future work.







Figure 6: Geographic distances of IP endpoints to IXP.

Traffic of the IXP

- Temporal: Strong diurnal behavior corresponding with business hours.
- Structural: Low rank - Implies alternative ways of data collection/measurement to avoid redundancy



(a) Daily pattern of top-10 tier-2 members.



(b) Traffic matrix energy.

Summary

By analyzing measurement data collected from many sources including the IXP, the paper presents a variety of interesting findings, inspiring further works in inspecting the special role of IXP in modern Internet. Findings include:

- Rich peering-peering links within the IXP
- A diversity of ASes in IXP that calls for better categorization methods instead of traditional tier-ed ISPs
- IXP traffic follows diurnal behavior as observed in ISP/Internet*

*: Quan *et al.*, "[When the Internet Sleeps: Correlating Diurnal Networks With External Factors](#)", IMC '14, November 2014

Discussion

- Many misconceptions corrected
 - Tier-1 ISPs DO peer at IXPs
 - IXPs ARE used for transit
 - There are MORE P-P links than customer-provider links
- IXP-external measurements
 - Why are they incomplete/misleading?
 - How could they be improved?
- Get hands dirty
 - Manually categorized ~400 ASes
- Reduce from many to few
 - Start looking for clues/features for as many as possible, then reduce the redundancy using mathematical methods