# On the Benefits of Using a Large IXP as an Internet Vantage Point

Nikolaos Chatzis, Georgios Smaragdakis\*, Jan Boettger, Thomas Krenc, Anja Feldmann TU Berlin/\*T-Labs

> Walter Willinger Niksun

#### Internet Vantage Points



Source: Arbor, "Internet Inter-Domain Traffic", SIGCOMM'10

#### A Unique Vantage Point: IXP



#### Internet eXchange Point (IXP)



## .. in reality IXP is more than a Switch



- Complex system
- A number of services are offered

For a survey: "There is More to IXPs than Meets they Eye", ACM SIGCOMM CCR, Oct. 2013

# Largest IXPs

Name		Main City	Members	Max Thr.	Av.Thr.	Traffic/day (ca. 2013)
	DE-CIX	Frankfurt	~500	2.5Tbps	I.4Tbps	~15 Petabytes/day
	AMS-IX	Amsterdam	~620	2.5Tbps	1.3Tbps	~14 Petabytes/day
	LINX	London	~440	1.5Tbps	ITbps	~11 Petabytes/day
	Equinix	All cities	~750	1.4Tbps	ITbps	~11 Petabytes/day
	DatalX	Moscow	~ 30	1.1Tbps	0.7Tbps	~7.5 Petabytes/day
	MSK-IX	Moscow	~600	ITbps	0.4Tbps	~4 Petabytes/day
	NetNod	Stockholm	~65	.5Tbps	0.3Tbps	~3 Petabytes/day

• ...

- Traffic comparable with this of Large Tier-I Networks:
  - AT&T: ~33 Petabytes/day (ca. July 2013)
  - Deutsche Telekom: ~16 Petabytes/day (ca. July 2013)

# Largest IXPs

Name		Main City	Members	Max Thr.	Av.Thr.	Traffic/day (ca. 2013)	
	DE-CIX	Frankfurt	~500	2.5Tbps	I.4Tbps	~15 Petabytes/day	
	AMS-IX	Amsterdam	~620	2.5Tbps	I.3Tbps	~14 Petabytes/day	
	LINX	London	~440	1.5Tbps	<b>I</b> Tbps	~11 Petabytes/day	
	Equinix	All cities	~750	1.4Tbps	ITbps	~11 Petabytes/day	
	DatalX	Moscow	~130	I.ITbps	0.7Tbps	~7.5 Petabytes/day	
	MSK-IX	Moscow	~600	1 Thee	0 4 These	ed Potabytos/day	
	NetNod	NetNod Stockholm ~65 Growth rates at the largest IXPs in Europe:					
	+ 10-20% new members/year						
	+ 50-100% more traffic/year						
	Traffic comparable with this + Offer 100Gbps ports						

- Traffic comparable with this o
  - AT&T: ~33 Petabytes/day (ca. July 2015)
  - Deutsche Telekom: ~16 Petabytes/day (ca. July 2013)

## Our Vantage Point: A Large IXP

#### Access to a Large European IXP (city metro)

Acknowledgments for the great collaboration with the IXP

#### What we know about this IXP from [1] in 2012 (traces from 2011)

- Detailed study of the "inside" picture of the IXP
- Main focus on connectivity
- Rich Ecosystem of IXP Members:
  - Access Networks
  - CDNs/Hosters
  - Transit Providers
  - Service Providers/Streamers
- Very dense peering among members, 50K+ out of the 78K possible, i.e., peering rate of 60%+

## **Open Questions**

What about the IXP as a vantage point for the Internet?

- Local vs. Global traffic visibility
- Stability vs. Trends in traffic flows
- What about the IXP as a vantage point for the commercial part of the Internet traffic?
  - Who is responsible for how traffic flows through the Internet: AS or Organizations or both?
  - What is the implication for traffic on peering links?

#### **IXP** Measurements

#### sFlow Data Collection:



I7 consecutive weeks of sFlow data,

#### weeks 35-51 2012 (August-December 2012)

- Sampling Rate: I/I6K packets
- Sampling Size: First 128 bytes of Ethernet frame
  - > 74 bytes of TCP payload
  - 86 bytes of UDP payload
- Traffic Volume Statistics:
  - Beginning: 443 members, ~12 Petabytes/day
  - End: 457 members, ~14 Petabytes/day

		1 week in Nov (week 45)	educated guesses of ground-truth
50	IPs	232,460,635	unknown $< 2^{32}$
fic	#ASes	42,825	approx. 43K
Lat e	Subnets	445,051	450K+
Η	countries	242	250

In a single week, we monitor traffic from essentially:

- all active ASes (recall there there are ~480 member ASNs, or 1% of all active ASes)
- all actively routed prefixes
- all countries

## IXP Server Visibility

- Servers are the engines of e-commerce and applications
- Server Identification:
  - Via HTTP:
    - String matching applied to the first line in response/request packets (e.g., GET, HEAD, POST, HTTP/1.{0,1}).
    - Commonly used HTTP header fields according to RFCs and W3C specifications.
  - Via HTTPS:
    - Step I: Consider IPs that use TCP port 443.
    - Step 2: Crawl each of these IPs for X.509 certificate chain.
    - Step 3: Check if the X.509 is valid.

#### Limitations:

- String matching may miss servers if there is no sufficient information in the payload.
- Some servers may mis-classified as clients when they behave as clients when communicating with other servers.
- HTTPS servers that do not use 443 will be missed.

### IXP Server Visibility

		1 week in Nov	educated guesses
		(week 45)	of ground-truth
	IPs	1,488,286	unknown
ver ffic	#ASes	19,824	unknown
laf	Subnets	75,841	unknown
L S	Countries	200	250

Traffic from:

- I 7% of the actively routed prefixes,
- ▶ 50% of the active ASes,
- > 200 of the countries

**Observations:** 

- Most popular ports: 80/8080 (80%), 1953 (~5%), 453 (~5%)
- ~250K HTTPS server IPs
- Many servers use multiple ports

#### IXP-external Traffic?



	Member AS	Distance 1	Distance $> 1$
			)
DO IPs	42.3%	45.0%	12.7%
. ∰ ∰ Prefixes	10.1%	34.1%	55.8%
a E ASes	1.0%	48.0%	50.1%
Traffic	67.3%	28.4%	4.3%
IPs	52.9%	41.2%	5.9%
ାତ୍ର Prefixes	17.2%	61.9%	20.9%
a ASes	2.2%	61.5%	36.3%
Traffic	82.6%	17.35%	0.05%

- Great visibility of non-IXP members: peer of IXP members, and peer of peers of IXP members!
- Beyond local traffic: 28% of total traffic and 17% of server traffic does not originate from an IXP member!

#### IXP: Local yet Global Visibility



## IXP: Local yet Global Visibility

		All IPs	Server IPs	All IPs	Server IPs
ra	nk	Country	Country	Network	Network
	1	US	DE	Chinanet	Akamai
	2	DE	US	Vodafone/DE	1&1
	3	CN	RU	Free SAS	OVH
	4	RU	FR	Turk Telekom	Softlayer
S	5	IT	GB	Telecom Italia	ThePlanet
Ħ	6	FR	CN	Liberty Global	Chinanet
	7	GB	NL	Vodafone/IT	HostEurope
	8	TR	CZ	Comnet	Strato
	9	UA	IT	Virgin Media	Webazilla
	10	JP	UA	Telefonica/DE	Plusserver

## IXP: Local yet Global Visibility

		All IPs	Server IPs	All IPs	Server IPs
rai	nk	Country	Country	Network	Network
	1	US	DE	Chinanet	Akamai
	2	DE	US	Vodafone/DE	1&1
	3	CN	RU	Free SAS	OVH
	4	RU	FR	Turk Telekom	Softlayer
S	5	IT	GB	Telecom Italia	ThePlaner
	6	FR	CN	Liberty Global	Chinanet
	7	GB	NL	Vodafone/IT	HostEurope
	8	TR	CZ	Comnet	Strato
	9	UA	IT	Virgin Media	Webazilla
	10	JP	UA	Telefonica/DE	Plusserver

### IXP: Local yet Global Traffic



### IXP Server Blind Spots

Which servers we can not see in the IXP and Why?

#### Source I: Large European Tier-1 ISP

- Full packet traces, thus very high accuracy in identifying servers and new URIs.
- Source II:Top-IM Alexa
  - Additional URIs from these retrieved from the IXP.
- Source III: Open DNS Resolvers
  - > 25K open resolvers in 12K ASes. We resolved all the URIs.

## IXP Server Blind Spots

- By combining all the IXP-external measurements we identified 600K server IPs, from which only 240K are new.
- The identity of the 240K "hidden" server IPs:
  - Private clusters of CDNs and Datacenters that are serving only customers of the same AS.
  - CDN servers in distant regions; This is to be expected as CDNs can well localize the content.
  - Traffic exchanged via private peering.
  - Hybrid Server Architectures if they are not using HTTP/ HTTPS.

#### Server Activity: Stable yet Changing



~70% of the total IXP traffic is due to server activity
~55% of the total IXP traffic is due to "stable" server IPs.

#### Server Activity: Global Observer



Steady increase of HTTPS traffic from 5% to 6% of total traffic

D

## Server Activity: Local Observer

- Deployment of New Servers and Business Trends:
  - Amazon EC2 in Europe: Increase of number of IPs last weeks of the year/before Christmas (e-commerce hot period).
  - First Installations of Netflix in Europe.

•

- New installation of Google caches within European ISPs.
- A number of outages of cloud providers with infrastructures located in Europe.
- IXP Resellers: Significant increase of traffic, the number of servers using resellers to send traffic doubled.

## **Open Questions**

What about the IXP as a vantage point for the Internet?

- Local vs. Global traffic visibility
- Stability vs. Trends in traffic flows



- What about the IXP as a vantage point for the commercial part of the Internet traffic?
  - Who is responsible for how traffic flows through the Internet: AS or Organizations or both?
  - What is the implication for traffic on peering links?

## Moving Beyond the AS-level View

- Given that a Large IXP is a unique vantage point, how we can use it to understand traffic flow in the Internet?
- What is the right abstraction?
  - ASes
  - Organizations, e.g., Google, Akamai, etc



# An Alternative Grouping of Server IPs

- We rely on recent results on how to map server IPs to commercial entities (organizations). See, e.g.,
  - DNS to Rescue: Discerning Content and Services in a Tangled Web, IMC'12.
  - Web Content Cartography, IMC'11.
  - Flexible Traffic and Host Profiling via DNS Randevouz, SATIN'II.
- For each server IP, we collect the following information from passive and active measurements:
  - Passive: URI
  - Active: related DNS queries/answers,
    - reverse DNS (hostname),
    - X.509 certificate (when available),

#### AS Heterogeneity: #Server IPs per Organization



- I 43 clusters with more than 1000 servers
- 6K clusters with more than 10 servers

### AS Heterogeneity: #Organizations per AS



#### AS-link Heterogeneity



### AS-link Heterogeneity



Akamai member AS peers with more than 400 networks.

### AS-link Heterogeneity



- Akamai member AS peers with more than 400 networks.
- Around 11% of the Akamai traffic does not traverse the Akamai link.



- Akamai member AS peers with more than 400 networks.
- Around 11% of the Akamai traffic does not traverse the Akamai link.
- Some networks do not receive traffic at all from the Akamai link.

Similar observations for other CDNs, e.g., CloudFlare

### Summary

- A large IXP is a single, well-localized vantage point with a great visibility of the Internet, not just their members.
- Having access to one of these large IXPs enables the tracking of new server deployments and trends in the Internet.
- Our study unveils significant heterogeneity of both ASes and AS-links.
- Our study challenges the mental model regarding the flow of Internet traffic.

\_\_\_\_\_

#### THANK YOU!