Traffic Asymmetry of IXPs: Characterization, Cause, and Impact

Leandro Bertholdo, Sandro Ferreira, Cesar Loureiro, João Ceron, Lisandro Granville, Ralph Holz, Roland van Rijswijk-Deij







UNIVERSITY OF TWENTE.

Why care about IXPs and asymmetry?

People connect at IXP to

- Reduces Latency
- Reduces Cost

Main benefit of IXPs

• Local transit stays local

The benefits are turning grey

- Open Peering in certain IXP can reach 38% of all Internet.
 - Is that good or bad?
- Some ASes on IXPs prefer to return traffic through transit links.
 - How many in terms of participants ?

Why care about IXPs and asymmetry?



Why is IXP asymmetry bad?



"OLD" Problems

- Wrong latency estimation
- Troubleshooting
- Optimization problem (we need ip geolocation with DNS resolution)

"NEW" Problems

- Small CDNs delivering traffic over lowquality paths
- Cost for CDNs (4x \$\$\$) -- Cloud services charging for intercontinental traffic.

What we know so far

Small CDNs are leaving IXPs and open peering, preferring to stay behind transit providers

Big CDNs prefer bi-lateral agreenments.

Some ASes do remote-peering on IXPs

IXPs are representative of the region! IX.br/SP covers 75% of LACNIC area, and 83% of Brazil

What we want to know

How many of those ASes really prefer using the IXP than the transit path?

How much asymmetric is the traffic on each IXP?

What can be done to improve?

Where? Which IXPs?

We applied on five IXPs

IXP	Rank	ASes	Open Peering	Traffic	Website
IX.br/SP	1	2,324	2,298	15 Tbps	ix.br
AMS-IX	3	847	571	11 Tbps	ams-ix.net
LINX	4	733	554	7 Tbps	linx.net
SIX	9	337	246	2 Tbps	seattleix.net
IX.br/RS	46	302	296	0.5 Tbps	ix.br

TABLE I: Selected IXPs by PeeringDB Ranking (May-2022)

Our challenge: How to measure?

Traceroutes (Ripe Atlas)

- Low coverage out-of-Europe (59% AMSIX and 4% ASes in IX.BR)
- Difficult problem: IXP transverse path identification

IXP data flows

- Hard to get access
- Not applicable everywere (Legislation)
- Just able to identify symmetry (sflow)

Routing dynamics

 IXP neighbors are stable – we limited to directed connected ASes

How did we measure?



https://anycast-testbed.nl/

How did we measure?



Our Results

(The IXP asymmetry big picture)

- **DEAF** NEIGHBORS: Some IXP neighbors IGNORE IXP routes
- **MUTE** NEIGHBORS: Some IXP neighbors FORWARD traffic to IXP but DO NOT do any prefix annouce
- The IXP path is being "depreferred" by IXP customers (ex. Prepend)



Details: IXP <u>network symmetry</u> with equal or more specific <u>prefix</u> size





Takeaway: Some operators generate asymmetry intentionally, but more than **half** we have consulted acknowledged **configuration mistakes**.

How about ASes peering on IXPs?

ixp	ASes	Unk	Symm	Hybrid	Ingress	Egress
AMS-IX	472	28	86.0 % 1	12	20	30
LINX	439	32	83.8 % 👃	10	22	35
IX.br/RS	220	18	94.1 % 🕇	2	3	7
SIX	204	22	84.2 %	9	12	26
IX.br/SP	1,879	261	90.7 %	13	20	116

Ouch: Do we have **DEAF** and **MUTE** neighbors? **Deaf**: Annouce prefix to IXP but Ignore IXP prefixes (egress-only) **Mute**: Return traffic on IXP but do not annouce any prefix (ingress-only)

Takeaway-1: In most cases, few ASes are responsible for asymmetry on IXPs.

Takeaway-2: Deaf and mute neighbors may be linked to **configuration mistakes** or a routing policy that prefers to use the IXP infrastructure as a backup path.

Details: Impact of more specific prefix on RTT



Takeaway: The use of unbalance prefix between IXP/ISP are prone to attract routes with higher RTT.

Is there any link between business type and asymmetry?

Takeaway: **ISPs are more symmetric** than expected. **Mobile operators are the most asymmetrical** and have good room for improvement on the IXPs we analyzed.

We also analyzed other CDNs They have several asymmetric prefixes (only-ingress)

CDNs sometimes deliver traffic from prefixes not annouced on IXPs (ex. Akamai)

IXP routing tables: only-egress and poor paths

Let's compare IXP routing table and global routing table prepends

Global routing table have around 10% (as3333 – RIPE view from RIPE-RIS)

Who is prepending at IXPs? (LINKS case)

Takeaway: We find IXP customers depreferring IXP routes when comparing with transit paths.

How about origin prepend? The impact of (as6939)

Takeway: long paths normally indicate poor quality routes. CDNs without quality-aware routing should de-peer with global networks in the IXP open peering model.

Conclusions in numbers

- Up to **24%** of ASes **avoid exchange** traffic over the IXP.
- 28% of IXPs paths are prepended
 - 15% IXP-customer over its clients
- Up to 8% of ASes filter out IXP routes.
- Up to 34% of IXP prefixes will not send traffic back

Possible solutions

- Informational: Our technique to assure IXP neighbors' symmetry can be used to build IXPmetrics (coverage, preference, asymmetry)
- **Business model:** IXPs can use local flow data to identify symmetrical paths and label with BGP Communities, **improving open peering**.
- **Standardization**: Anycast networks and CDNs demand a special treatment from routing peers.
 - Special AS-Range for CDNs and anycast networks (eg. DNS authoritative providers)
 - New IXP communities
 - draft-wilhelm-grow-anycast-community-01, Jul. 2022.

More information

https://github.com/LMBertholdo/ixp-symmetry-rate

THANK YOU!

QUESTIONS?

Leandro M. Bertholdo <leandro.bertholdo@gmail.com>

TWENTE.

INSTITUTO FEDERAL Rio Grande do Sul