

# On the Asymmetry of Internet eXchanges Points. Why IXPs and CDNs should care?

Leandro M. Bertholdo\*, Sandro L. A. Ferreira<sup>§</sup>, João M. Ceron<sup>†</sup>,  
Lisandro Zambenedetti Granville<sup>‡</sup>, Ralph Holz\*, Roland van Rijswijk-Deij\*<sup>¶</sup>

\*University of Twente, Enschede, The Netherlands - {l.m.bertholdo, r.holz, r.m.vanrijswijk}@utwente.nl

<sup>§</sup>Federal Institute of Technology, Porto Alegre, Brazil - sandro@ifrs.edu.br

<sup>†</sup>Botlog.org, Brazil - ceron@botlog.org

<sup>‡</sup>Federal University of Rio Grande do Sul, Porto Alegre, Brazil - granville@inf.ufrgs.br

<sup>¶</sup>NLnet Labs, Amsterdam, The Netherlands

# What is the problem?

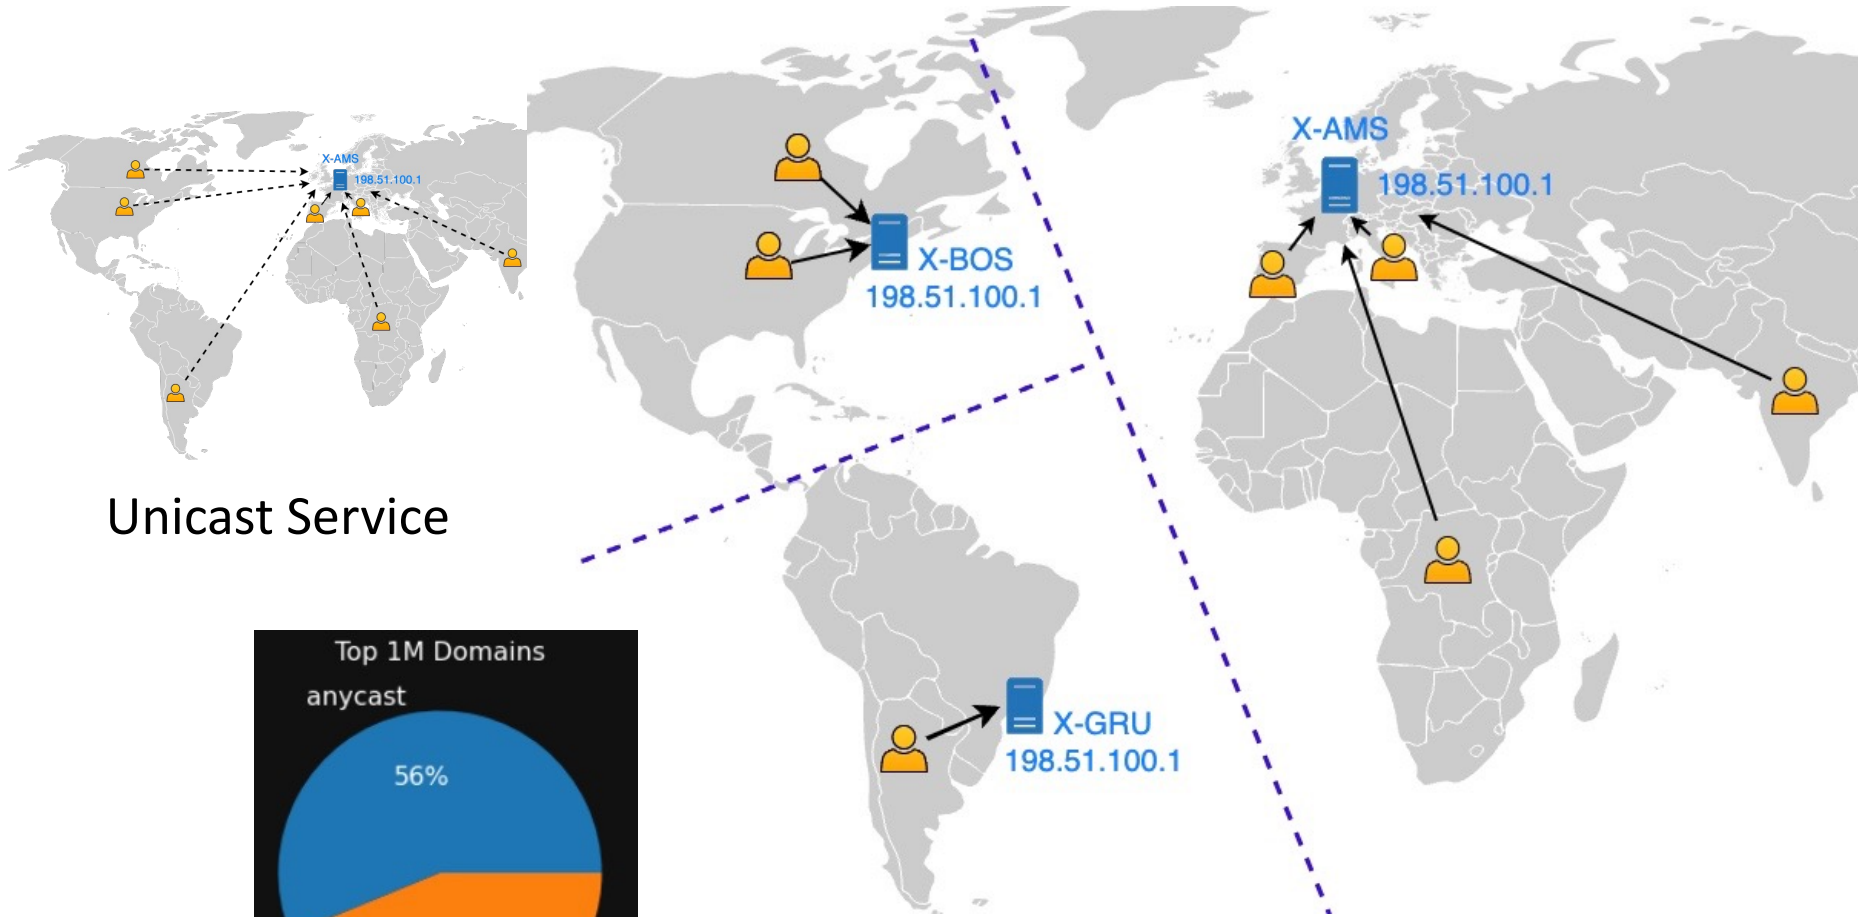


Small CDNs are  
disconnecting from IXPs

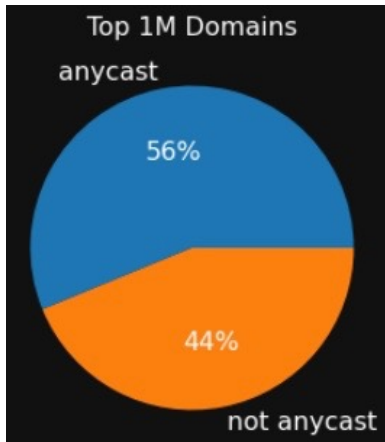
Small content providers are  
disconnecting from IXPs to connect  
to major transit providers

\* Small CDNs are heavily based on anycast

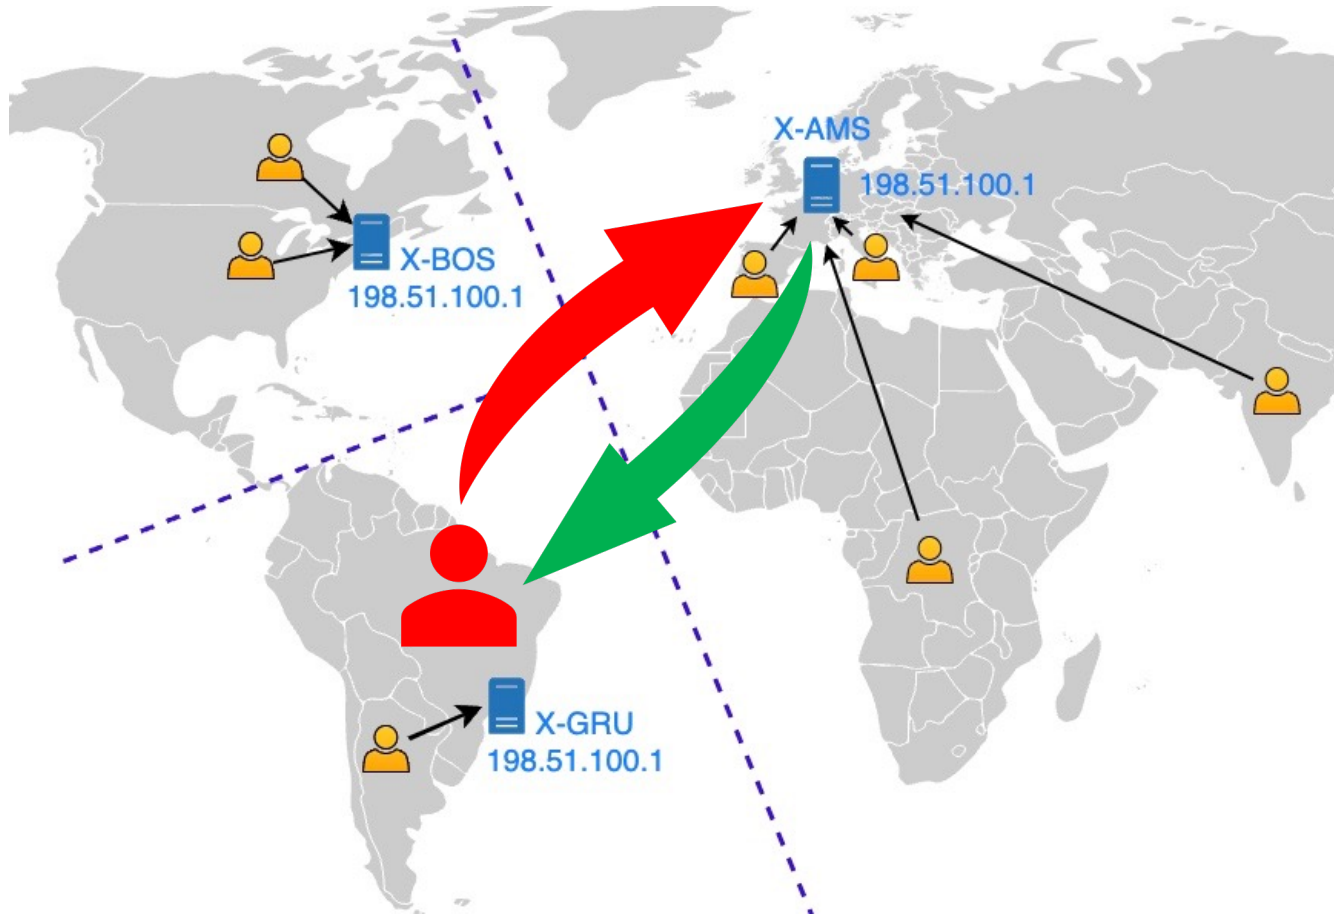
# First: How anycast works?



Unicast Service




First: How anycast works?

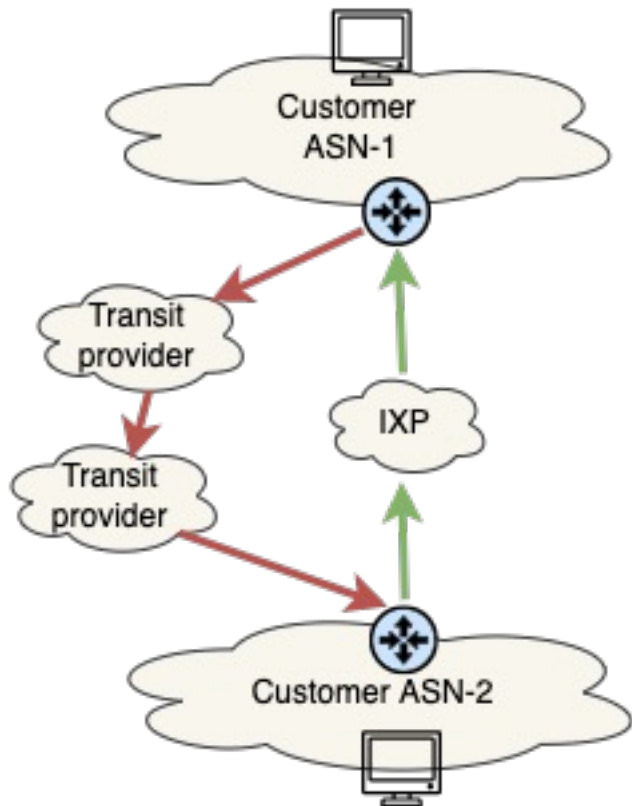


ISPs' Routing Policies

# Why IXPs ?



# Why is routing asymmetry bad?



## OLD Problems

- Wrong latency estimation
- Troubleshooting
- Optimization problem

## “NEW” Problems

- Low quality paths:
  - Affect CDNs
- Cloud services charge intercontinental traffic
  - Cost ( 4x \$\$\$\$ )

What we  
want to  
know

---

How many ASes prefer  
using the IXP than the  
transit path?

---

How asymmetric is the  
traffic on each IXP?

---

What can be done to  
improve?

# 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

- Not applicable everywhere (Legislation)
- Just able to identify symmetry (sflow)

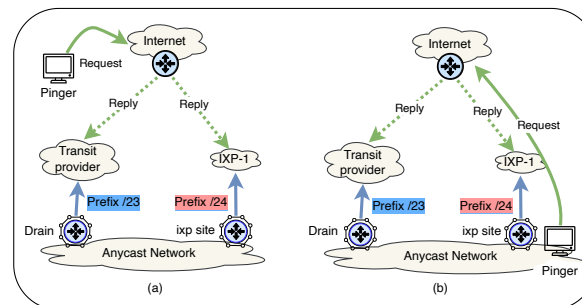
## Routing dynamics

- IXP neighbors are stable – we limited to directed connected ASes

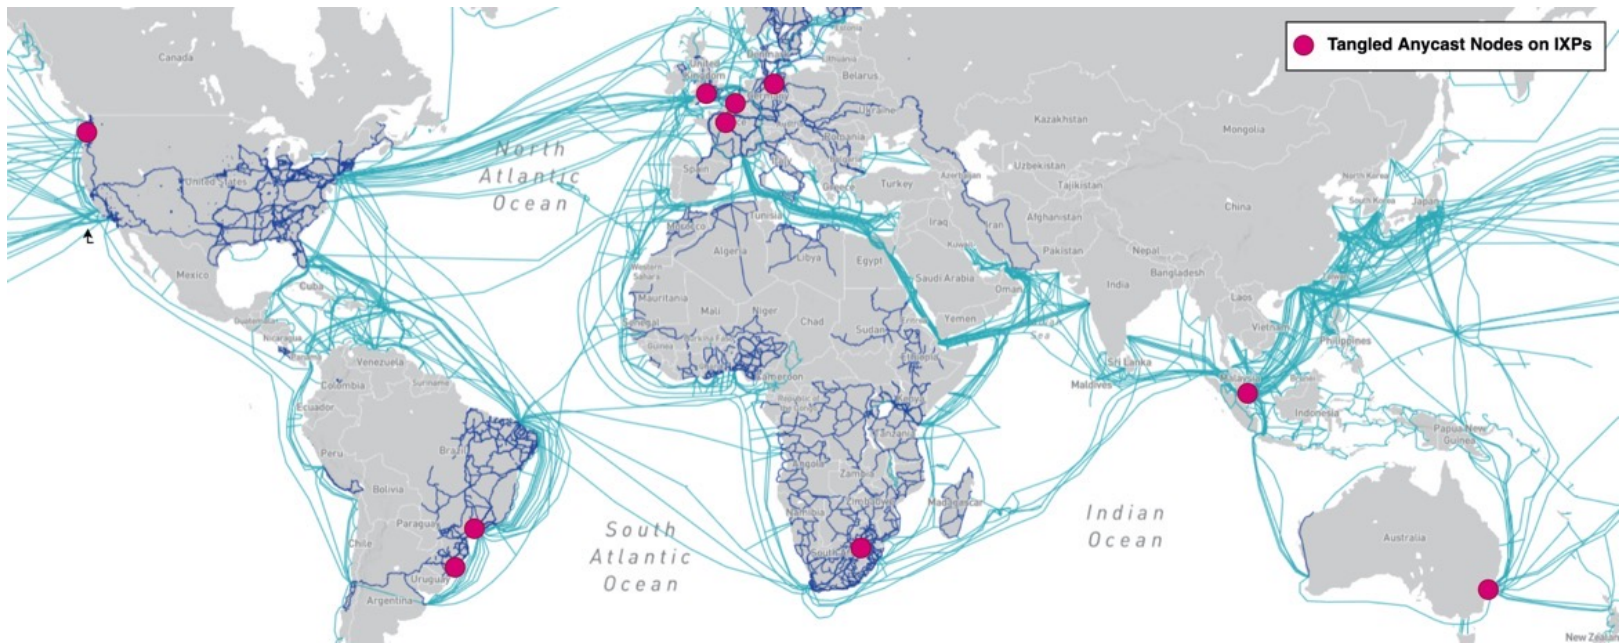


# What we did?

- We proposed a new method of **anycast active measurements**
- **We use “anycast as a measurement tool”**
- In this method we
  - Connected anycast sites on IXPs and one ISP
  - Manipulate routing
  - Actively generate traffic for **6 million /24 networks** on IXP and Transit
  - Map the behaviour of each individual ASe connected in each IXP
  - We map up to **89% of all ASNs** on the Internet in **15 minutes**.
  - Comparing with RIPE Atlas we map **91% of AMSIX ASes** and **79% SPO.IX.BR**



# How do we test? Anycast testbed (Tangled)



<https://anycast-testbed.nl/>

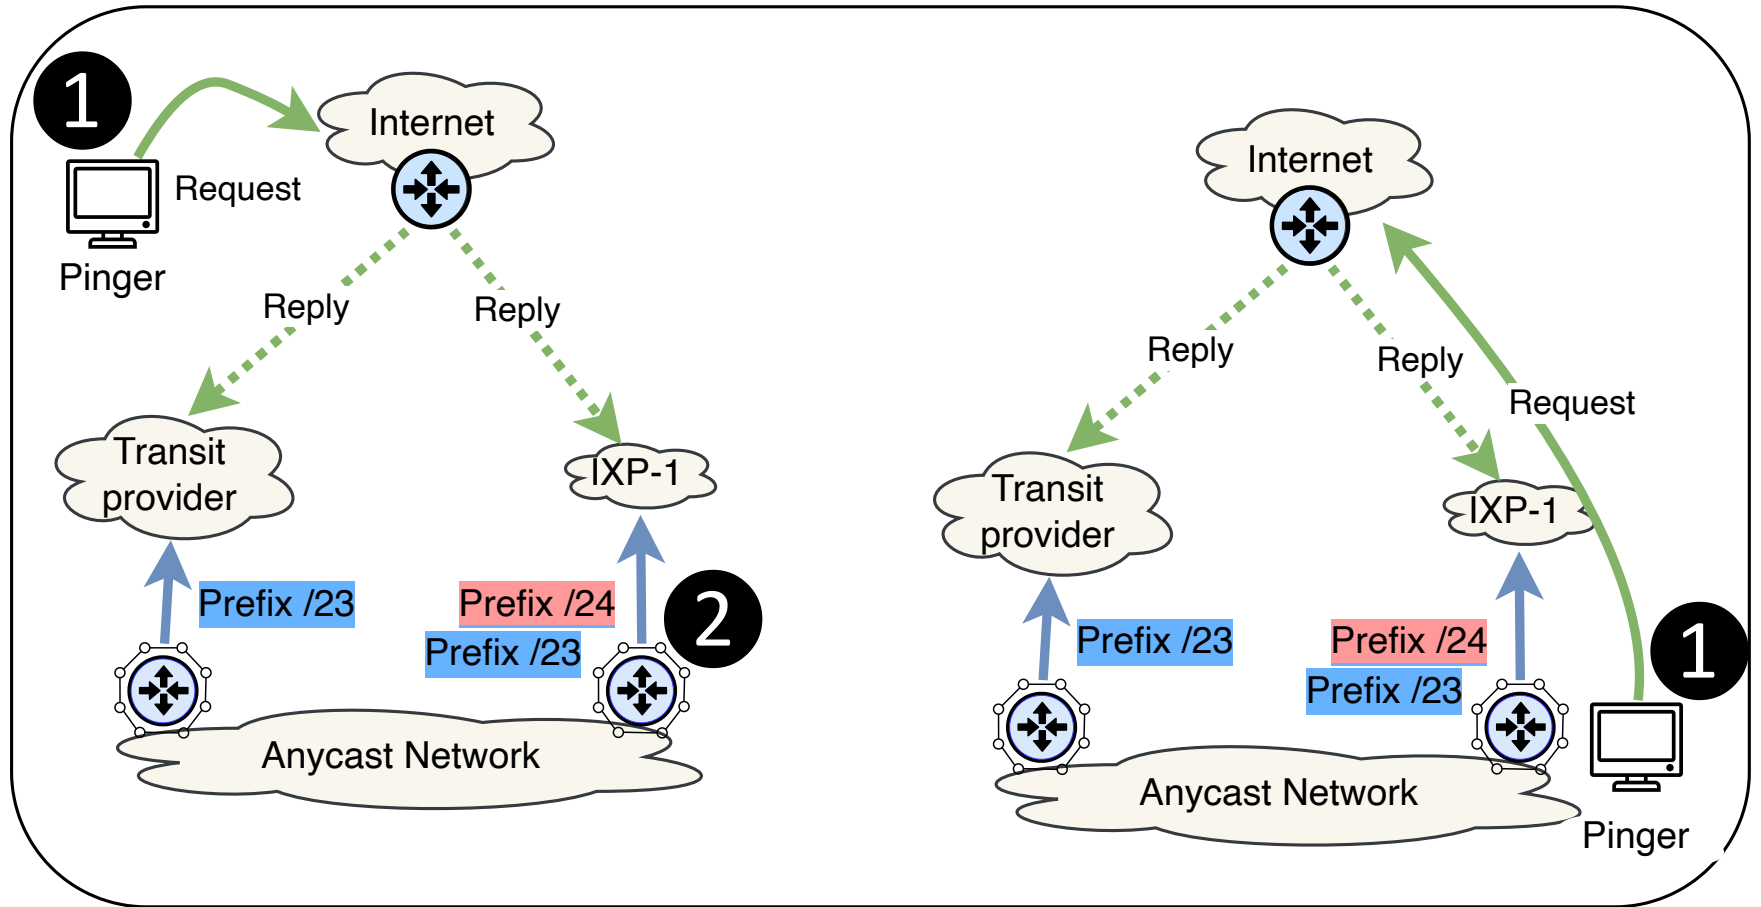
# How do we test? Which IXPs

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)

# How do we test? Methodology


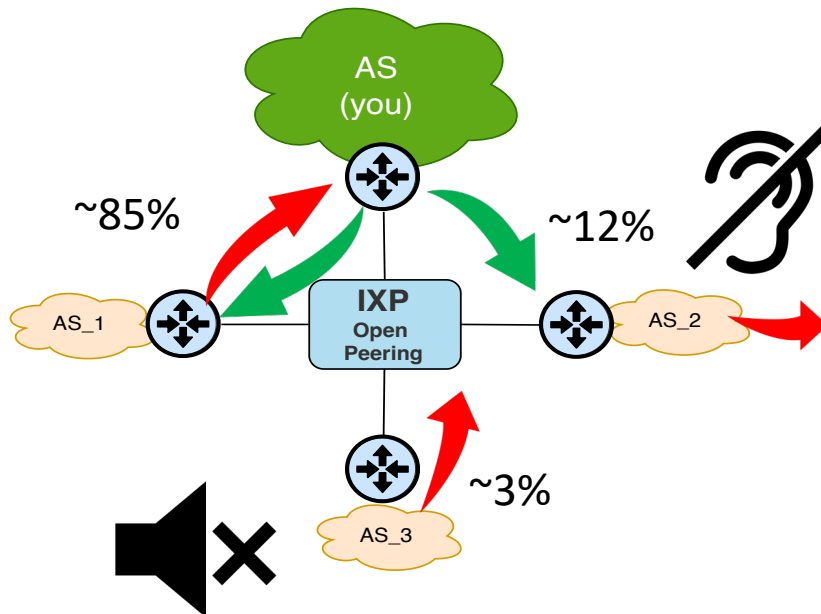


**1** Source

**2** Prefix Size

# What have we learned? (The IXP 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 announce
- **The IXP path is being “depreferred” by IXP customers**



# Details: IXP network symmetry with equal or more specific prefix size

/24

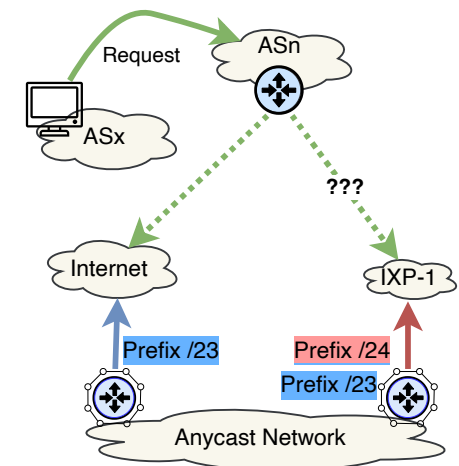
ixp	Neig. Net.	Symmetric	Only Ingress	Only Egress
AMS-IX	90,064	79.4%	6.6%	13.8%
LINX	66,040	88.5%	7.0%	4.2%
IX.br/RS	7,917	78.9%	1.1%	20.0%
SIX	31,286	88.1%	3.7%	8.1%
IX.br/SP	35,327	85.3%	1.7%	12.2%

TABLE II: Network symmetry using more specific prefix

/23


ixp	Neig. Net.	Symmetric	Only Ingress	Only Egress
AMS-IX	85,967	63.2% ↓↓	2.1% ↓↓	34.4% ↑↑
LINX	65,258	74.3% ↓↓	6.0% ↓	19.3% ↑↑
IX.br/RS	7,903	76.7% ↓	0.9% →	22.2% ↑
SIX	31,310	86.3% ↓	3.4% →	10.1% ↑
IX.br/SP	34,984	85.4% →	0.8% →	13.6% ↑

TABLE III: Network symmetry using same size prefix




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

# 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.



# Details: Mapping AS-Level behavior on IXPs

ixp	ASes	Unk	Symm	Hybrid	Ingress	Egress
AMS-IX	472	28	86.0 % ↑	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

TABLE IV: ASes symmetry using more specific prefix

Ouch: Do we have **DEAF** and **MUTE** neighbors?

**Deaf:** Announce prefix to IXP but Ignore IXP prefixes (egress-only)

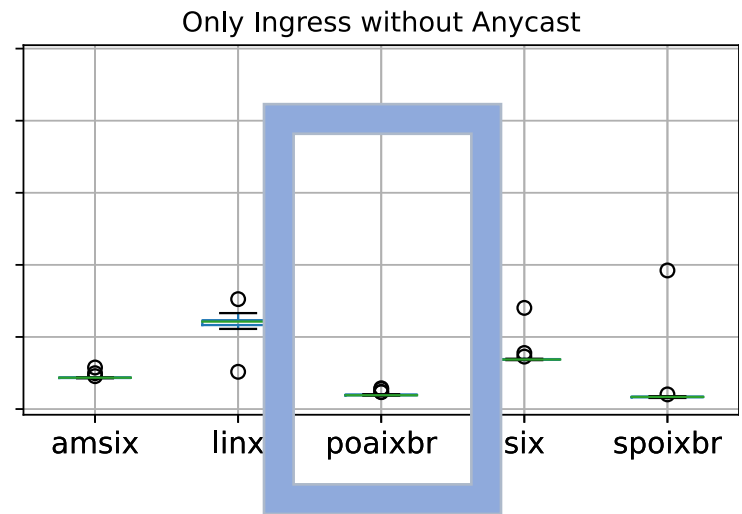
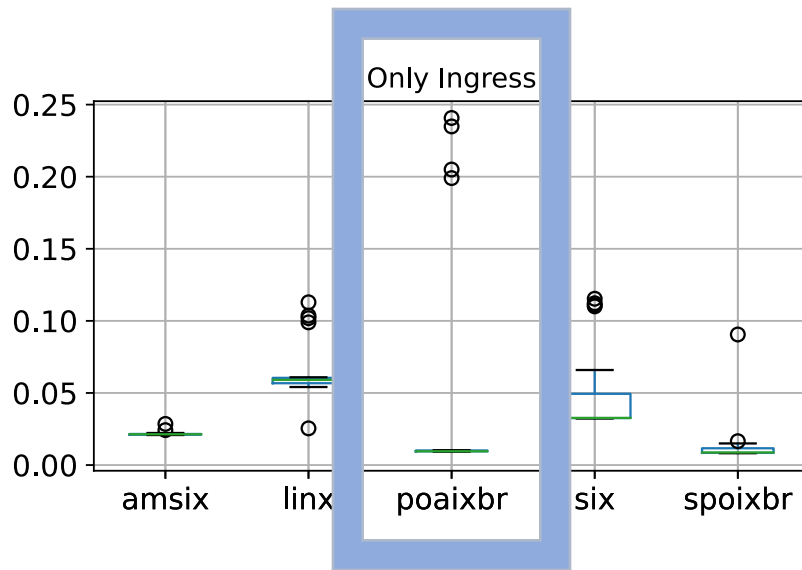
**Mute:** Return traffic on IXP but do not announce any prefix (ingress-only)

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

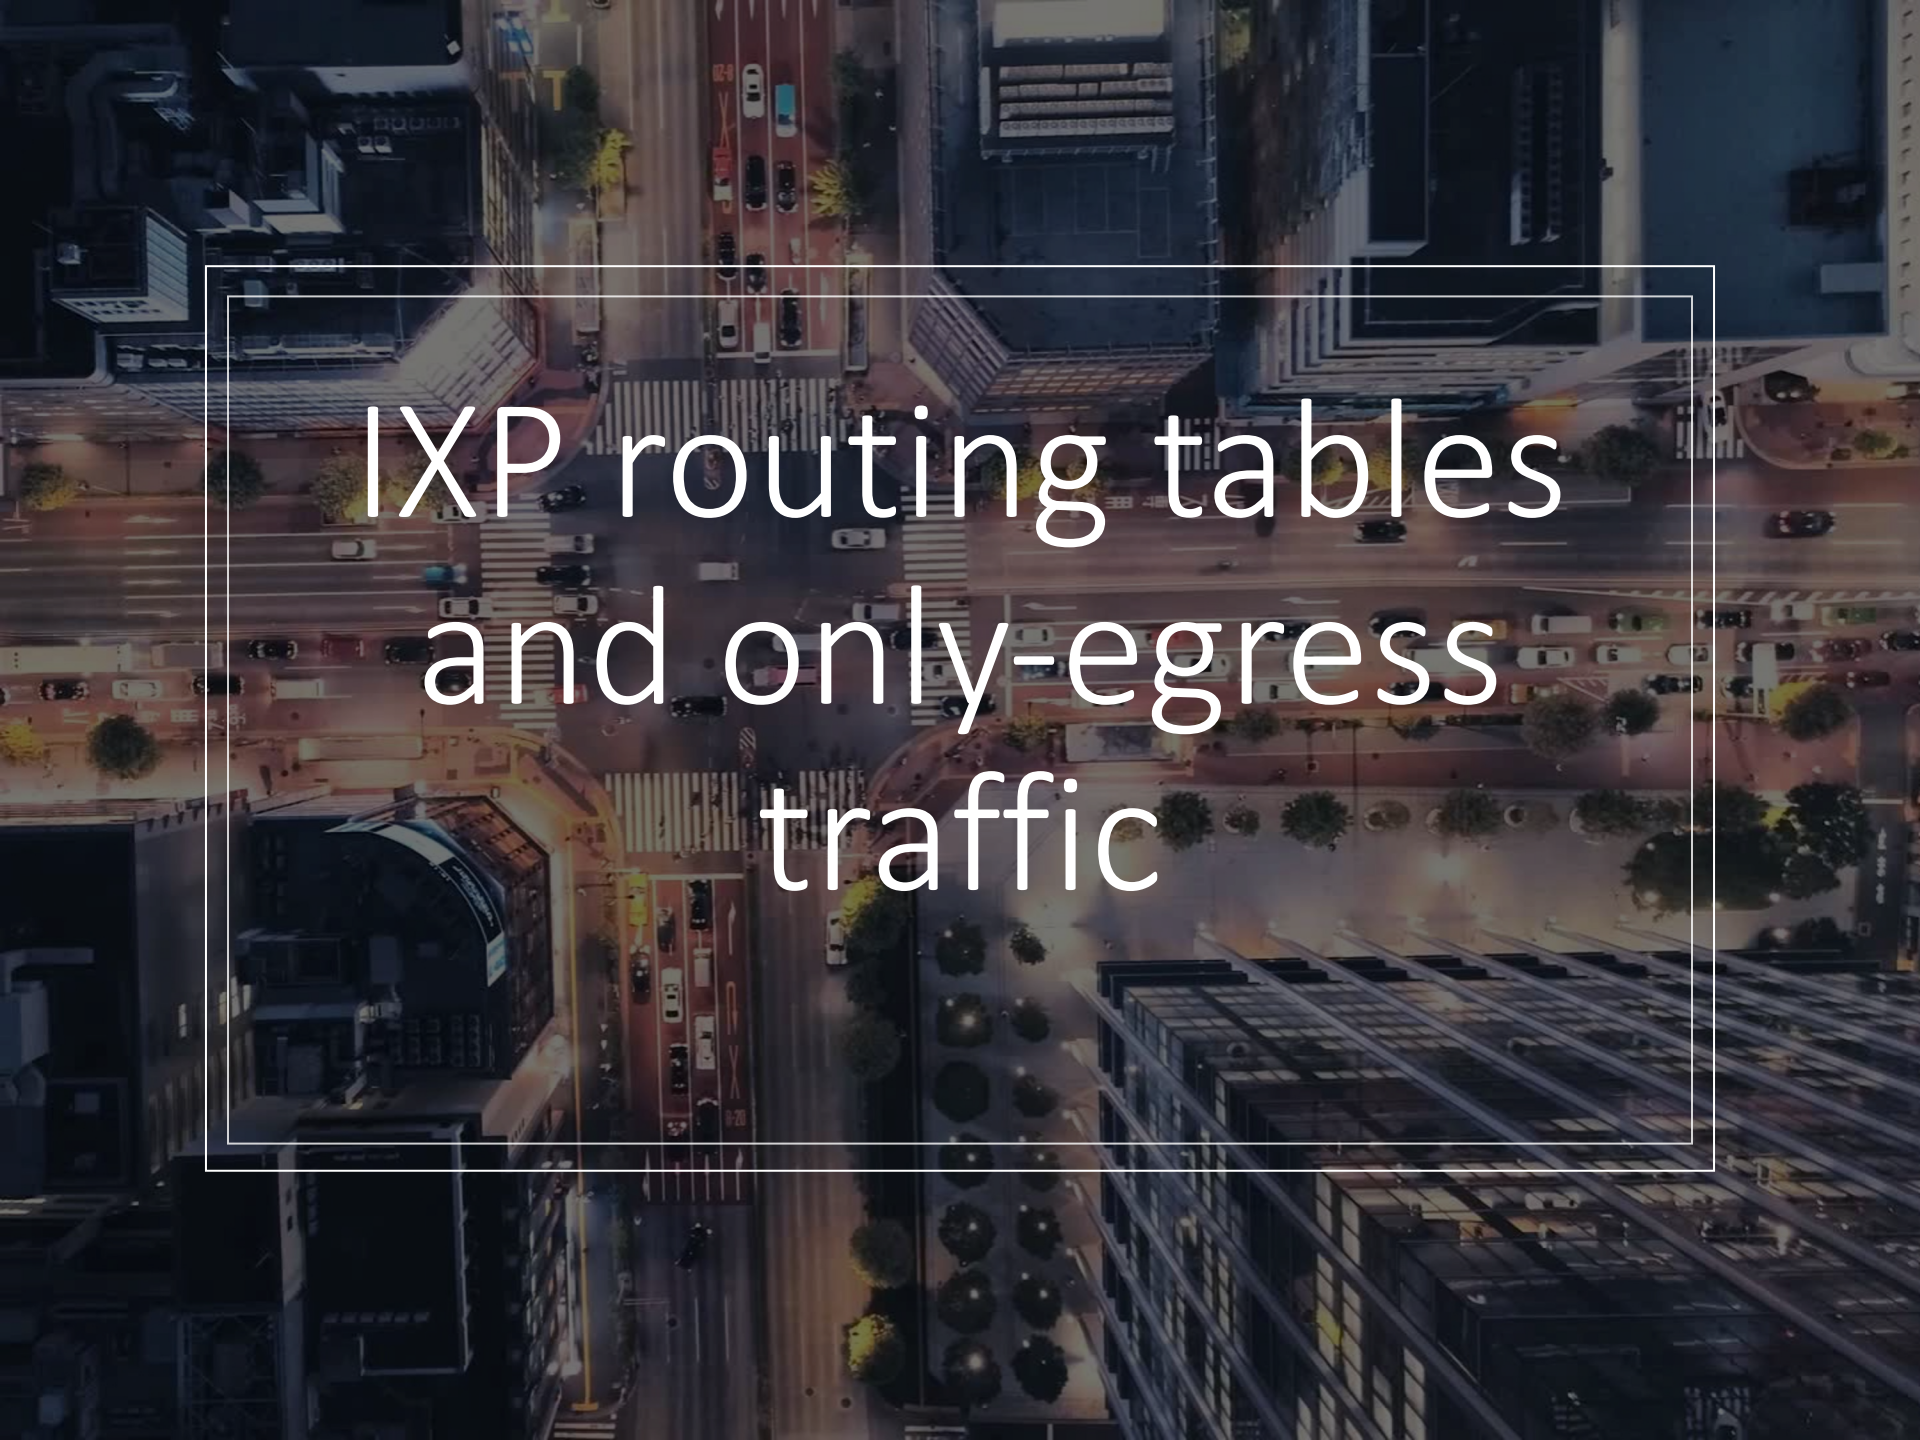
**Takeaway-2:** Deaf and mute neighbors may be linked to **configuration mistakes** or a routing policy that prefers to use the **IXP as a backup path**.

# We also analyzed other CDNs

They have several asymmetric prefixes (only-ingress)




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

An aerial, top-down view of a city street at night. The street is illuminated by streetlights, and several cars are visible on the road. Buildings with lit windows line the street. The overall scene is dark with some highlights from the lights.

# IXP routing tables and only-egress traffic


# Can depreferred paths be increasing ingress-only asymmetry? Let's look IXP routing tables...




LINX: ~30% all paths prependded

All IXPs have between 26-31% prependded paths

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

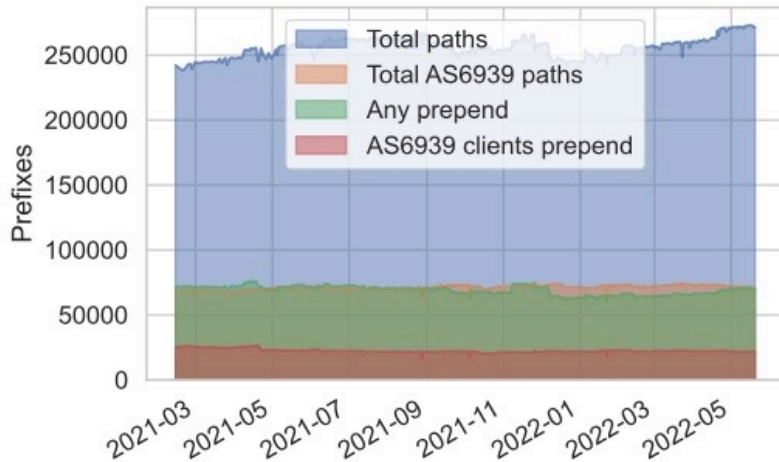


# Who is prepending at IXPs? (LINX case)

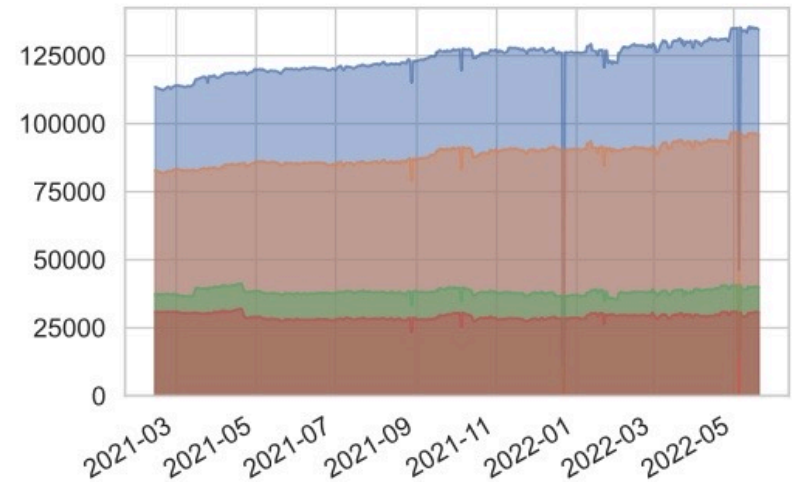


**Takeaway:** We find IXP customers deprefering IXP routes when comparing with transit paths.

# How about origin prepend? The impact of (as6939)



(a) AMSIX.



(b) SIX.

**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

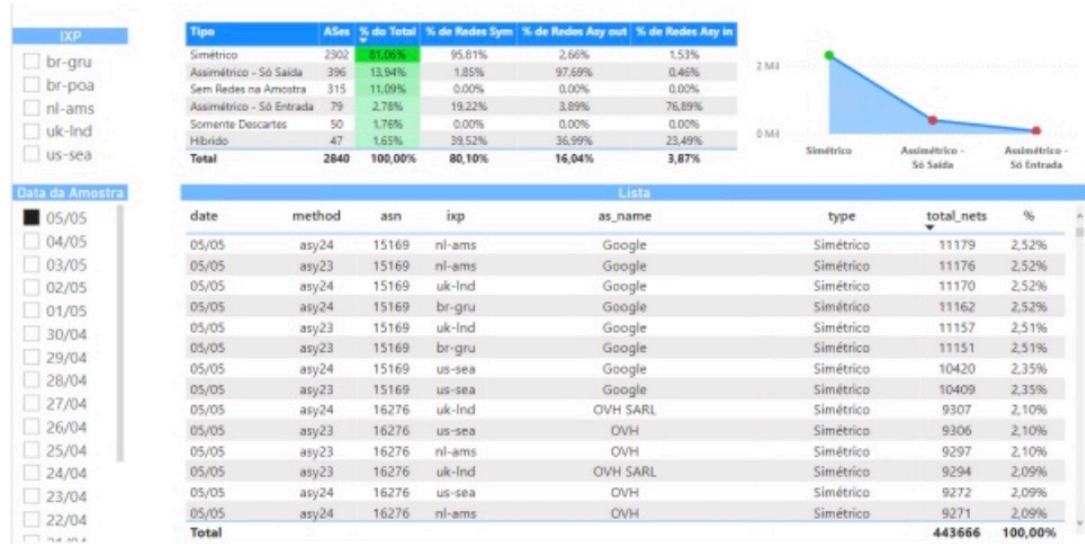
---

- **Informational:** This method can be used at scale to build IXP metrics (coverage, preference, asymmetry).
- **Business model:** IXPs can use local flow data to identify symmetrical paths improving multi-lateral views.
- **Standardization:** Anycast networks demand a special treatment from routing peers
  - [draft-wilhelm-grow-anycast-community-01](#), Jul. 2022.
  - Special AS-Range
  - New protocols






More information



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

<https://paaddos.nl>

# THANK YOU!



## QUESTIONS?

Leandro M. Bertholdo <l.m.bertholdo@utwente.nl>  
Design and Analysis of Communication Systems (DACs)  
Electrical Engineering, Mathematics and Computer  
Science (EEMCS)



UNIVERSITY OF  
TWENTE.