lacnic 45

25 - 28 MAYO / CIUDAD DE PANAMÁ, PANAMÁ

Behind the Routes: Measuring BGP Instability on the Internet

Milena Bueno, Pedro Marcos

milenabueno@furg.br pbmarcos@furg.br

FURG

UNIVERSIDADE FEDERAL

DO RIO GRANDE

Traffic delivery is a fundamental aspect of Internet operations

Traffic delivery is a fundamental aspect of Internet operations

BGP is responsible for routing decisions between Autonomous Systems

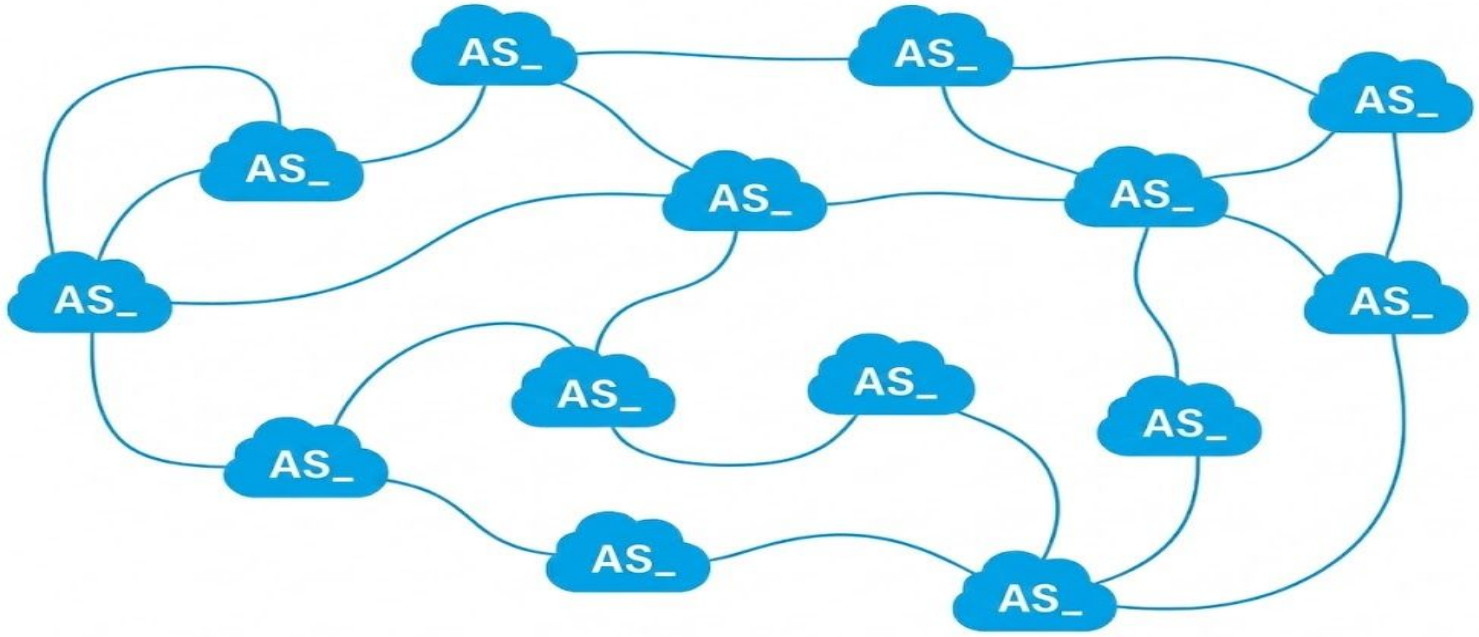

In practice, routes are never fully stable…

In practice, routes are never fully stable…

Traffic engineering is a constant need, driven by application demands, network failures, and customer satisfaction

In practice, routes are never fully stable…

Traffic engineering is a constant need, driven by application demands, network failures, and customer satisfaction

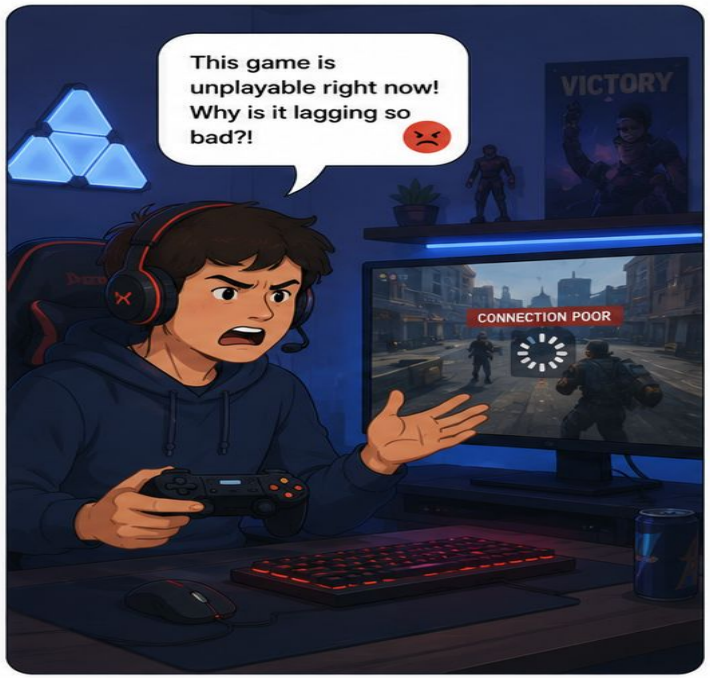

In practice, routes are never fully stable…

Traffic engineering is a constant need, driven by application demands, network failures, and customer satisfaction

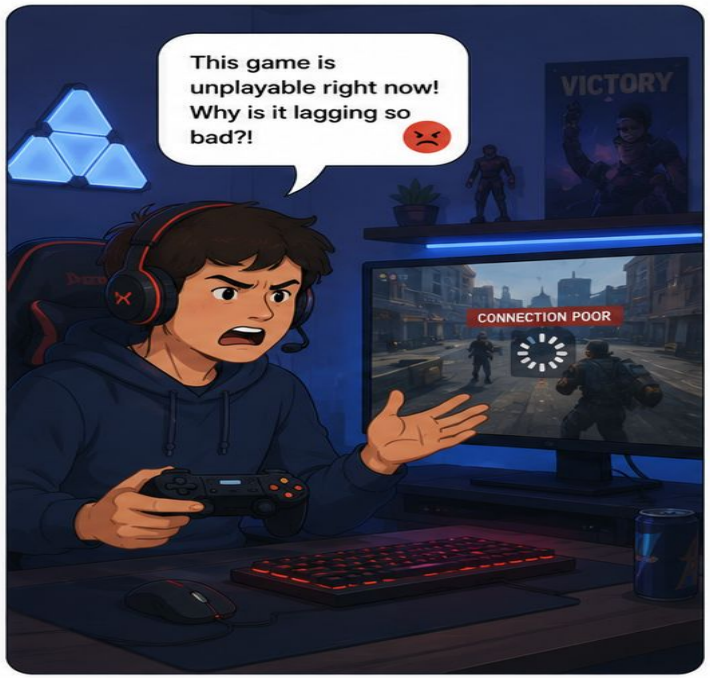

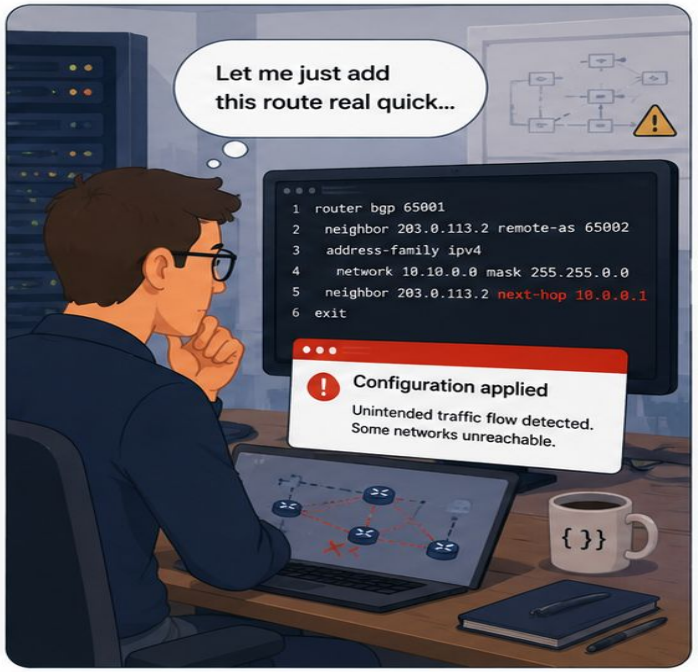

In practice, routes are never fully stable…

Traffic engineering is a constant need, driven by application demands, network failures, and customer satisfaction

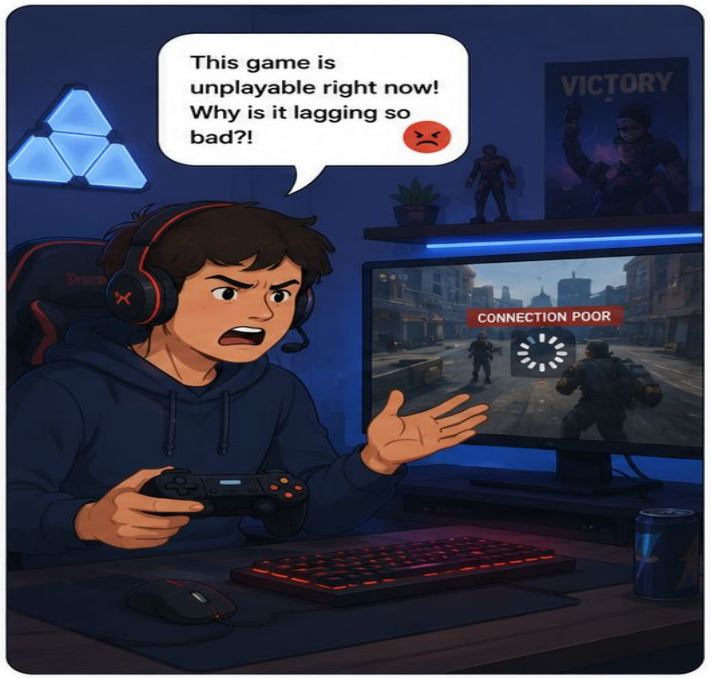

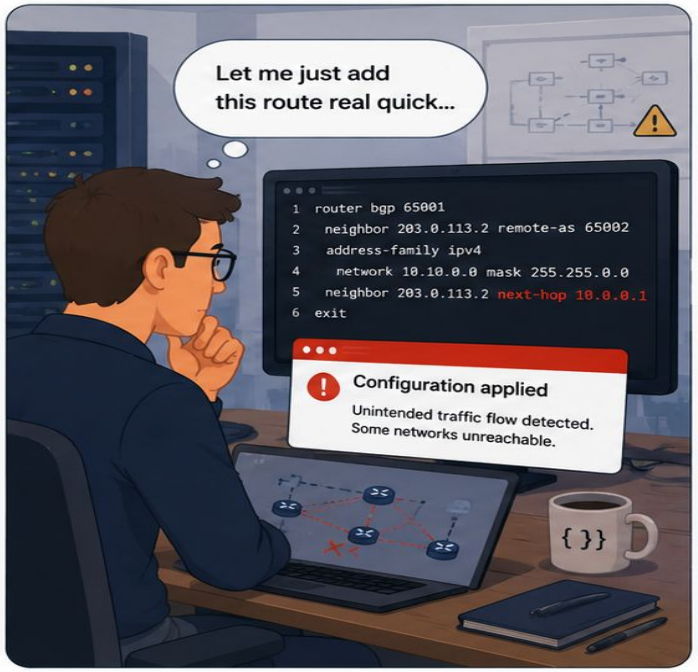

When instability occurs, it has consequences

When instability occurs, it has consequences

Increased Latency

When instability occurs, it has consequences

Increased Latency

Packet Loss

When instability occurs, it has consequences

Increased Latency

Packet Loss

Service Outages

Our goal is to understand the dynamics of BGP routing instability, investigating the following aspects:

Our goal is to understand the dynamics of BGP routing instability, investigating the following aspects:

Q1 - How unstable is the BGP?

Our goal is to understand the dynamics of BGP routing instability, investigating the following aspects:

Q1 - How unstable is the BGP? **Q2** - What is changing?

Our goal is to understand the dynamics of BGP routing instability, investigating the following aspects:

Q1 - How unstable is the BGP?

Q2 - What is changing?

Q3 - How do routing changes propagate across the Internet ?

lacnic

Data sources and methodology

Routing Data

lacnic

Data sources and methodology

Routing Data

RIBs

+

Updates

lacnic

Data sources and methodology

Routing Data

RIPE NCC
RIPE NETWORK COORDINATION CENTRE

RIBs

+

Updates

Selected Collectors

RRC00 - Amsterdam, NL

Global View **Multihop**

RRC24 - Montevideo, UY

LACNIC region **Multihop**

lacnic

Data filtering

Only monitors announcing a full routing table were considered

Full table thresholds

IPv4

≥ 850.000
announced prefixes

IPv6

≥ 180.000
announced prefixes

Data filtering

Only monitors announcing a full routing table were considered

Full table thresholds

IPv4

≥ 850.000
announced prefixes

IPv6

≥ 180.000
announced prefixes

Collection Period

April 2026

01 → 15

15 DAYS

lacnic

We investigate the instability across three scenarios

(Regional prefixes identified using LACNIC delegated files)

We investigate the instability across three scenarios

(Regional prefixes identified using LACNIC delegated files)

GLOBAL

IPv4 6.6M
IPv6 2.1M

Announcements

198M

Withdrawals

49 M

lacnic

We investigate the instability across three scenarios

(Regional prefixes identified using LACNIC delegated files)

GLOBAL

IPv4 6.6M
IPv6 2.1M

Announcements
198M

Withdrawals
49 M

LACNIC

IPv4 104k
IPv6 26.4k

Announcements
21.7M

Withdrawals
1.4 M

We investigate the instability across three scenarios

(Regional prefixes identified using LACNIC delegated files)

GLOBAL

IPv4 6.6M
IPv6 2.1M

Announcements

198M

Withdrawals

49 M

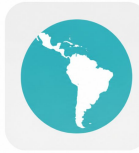

LACNIC

IPv4 104k
IPv6 26.4k

Announcements

21.7M

Withdrawals

1.4 M

PANAMA

IPv4 404
IPv6 58

Announcements

9.5k

Withdrawals

2.3k

Our goal is to understand the dynamics of BGP routing instability, investigating the following aspects:

Q1 - How unstable is the BGP?

Not all updates are equal but most of them matter …

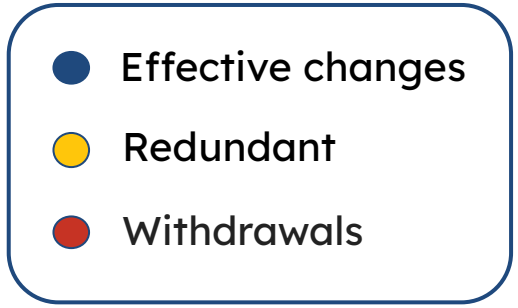

Not all updates are equal but most of them matter …

- Effective changes
- Redundant
- Withdrawals

GLOBAL

- 66.5% :

IPv4 51.9% | IPv6 48.1%

- 13.7%

IPv4 58.3% | IPv6 41.7%

- 19.8%

IPv4 58.3% | IPv6 41.7%

Not all updates are equal but most of them matter …

- Effective changes
- Redundant
- Withdrawals

GLOBAL

- 66.5% :

IPv4 51.9% | IPv6 48.1%

- 13.7%

IPv4 58.3% | IPv6 41.7%

- 19.8%

IPv4 58.3% | IPv6 41.7%

LACNIC

- 81.7%:

IPv4 21.1% | IPv6 78.9%

- 12.1%

IPv4 26.5% | IPv6 73.5%

- 6.1%

IPv4 27.3% | IPv6 72.7%

Not all updates are equal but most of them matter …

- Effective changes
- Redundant
- Withdrawals

GLOBAL

- 66.5% :

IPv4 51.9% | IPv6 48.1%

- 13.7%

IPv4 58.3% | IPv6 41.7%

- 19.8%

IPv4 58.3% | IPv6 41.7%

LACNIC

- 81.7%:

IPv4 21.1% | IPv6 78.9%

- 12.1%

IPv4 26.5% | IPv6 73.5%

- 6.1%

IPv4 27.3% | IPv6 72.7%

PANAMA

- 63.6%

IPv4 82.3% | IPv6 17.7%

- 16.9%

IPv4 83.2% | IPv6 16.8%

- 19.5%

IPv4 97.5% | IPv6 2.5%

But what is actually changing inside these updates?

lacnic

But what is actually changing inside these updates?

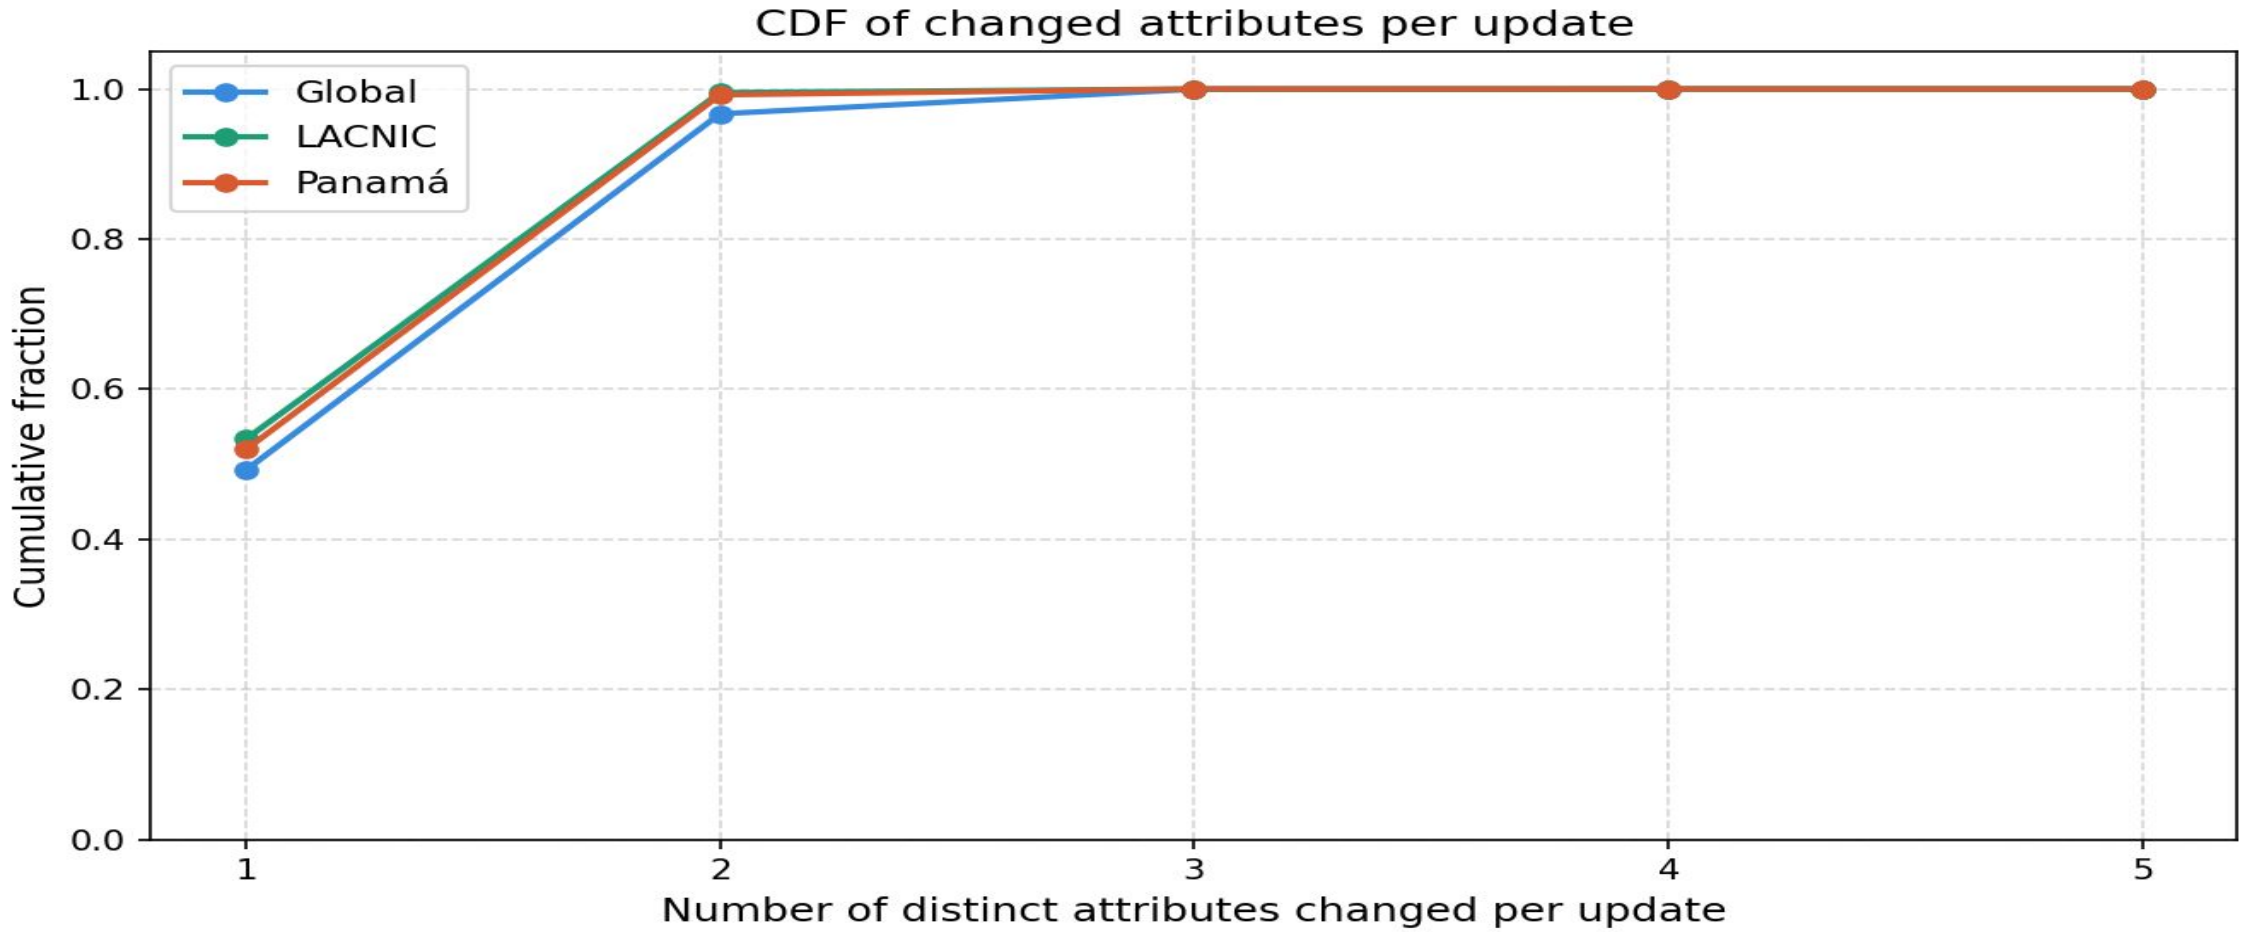

But what is actually changing inside these updates?

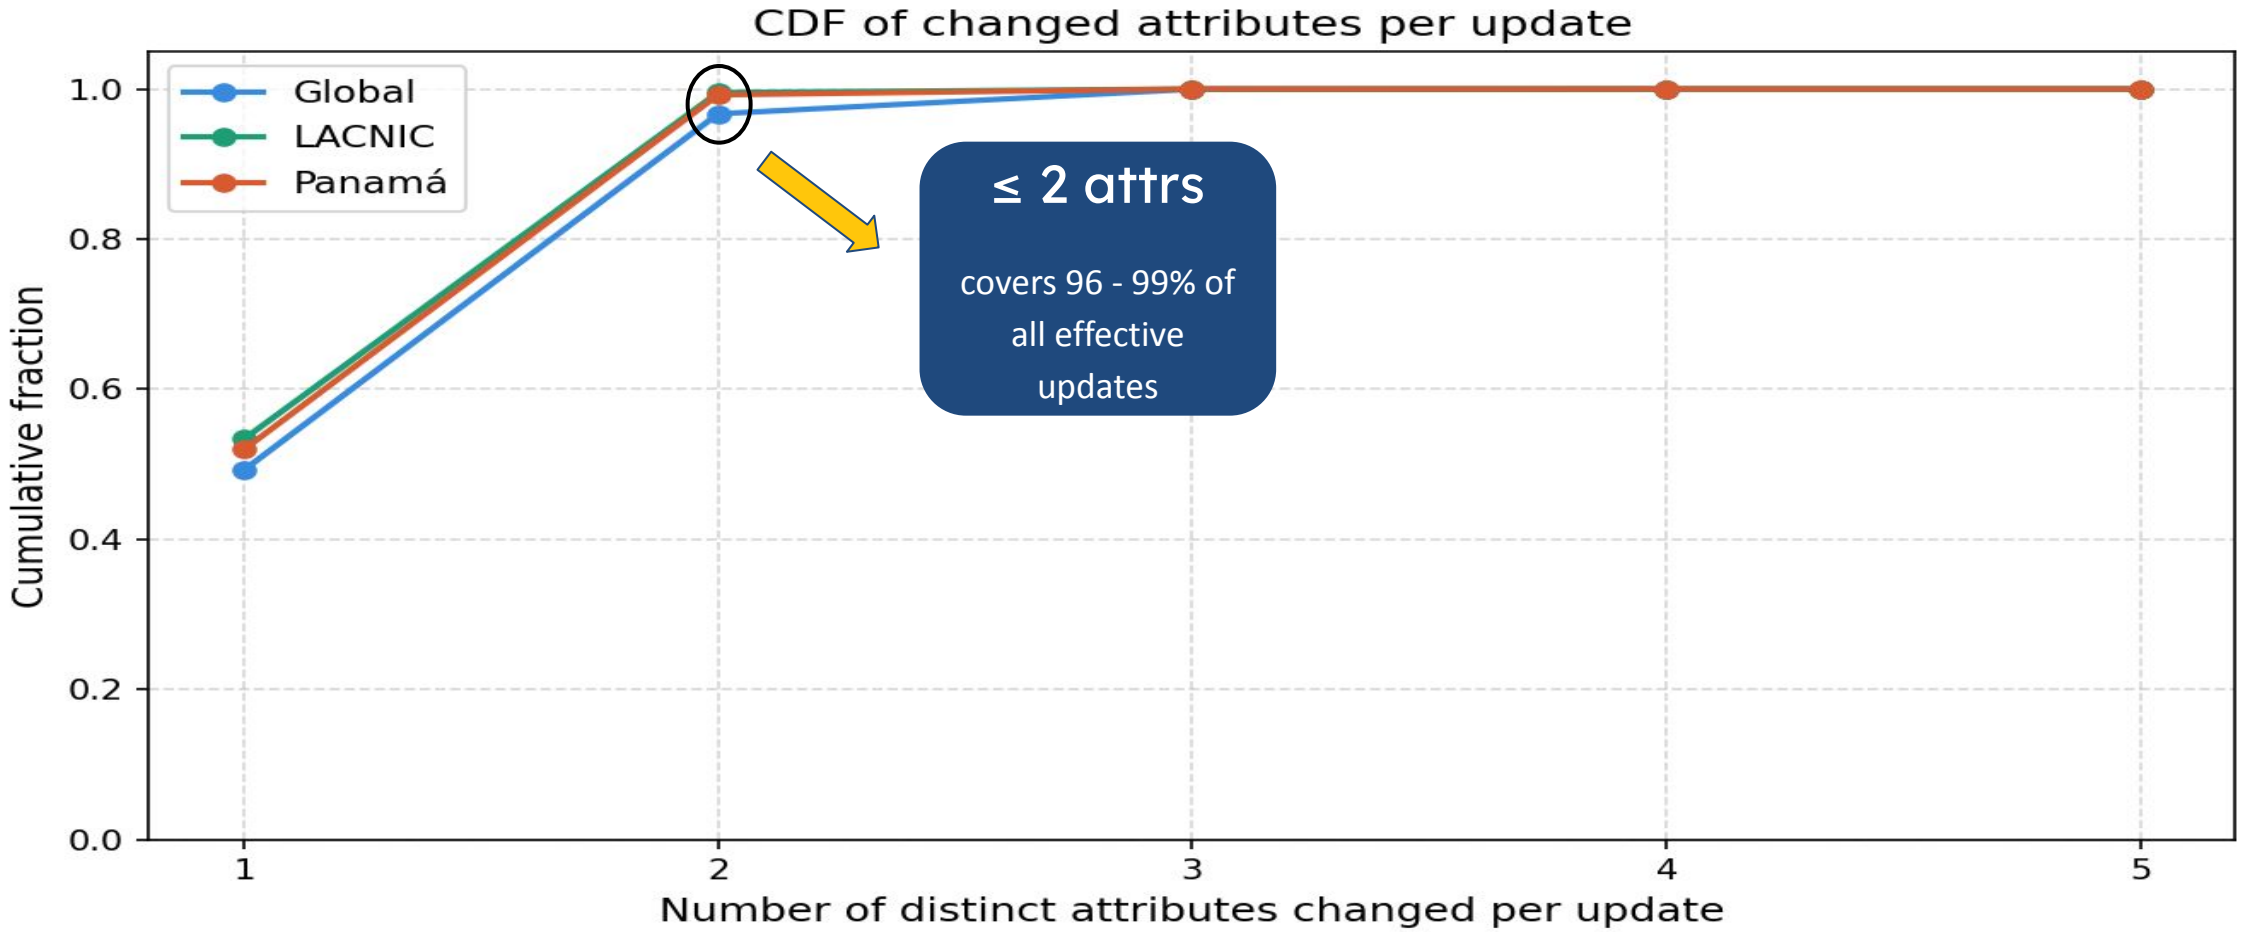

But what is actually changing inside these updates?

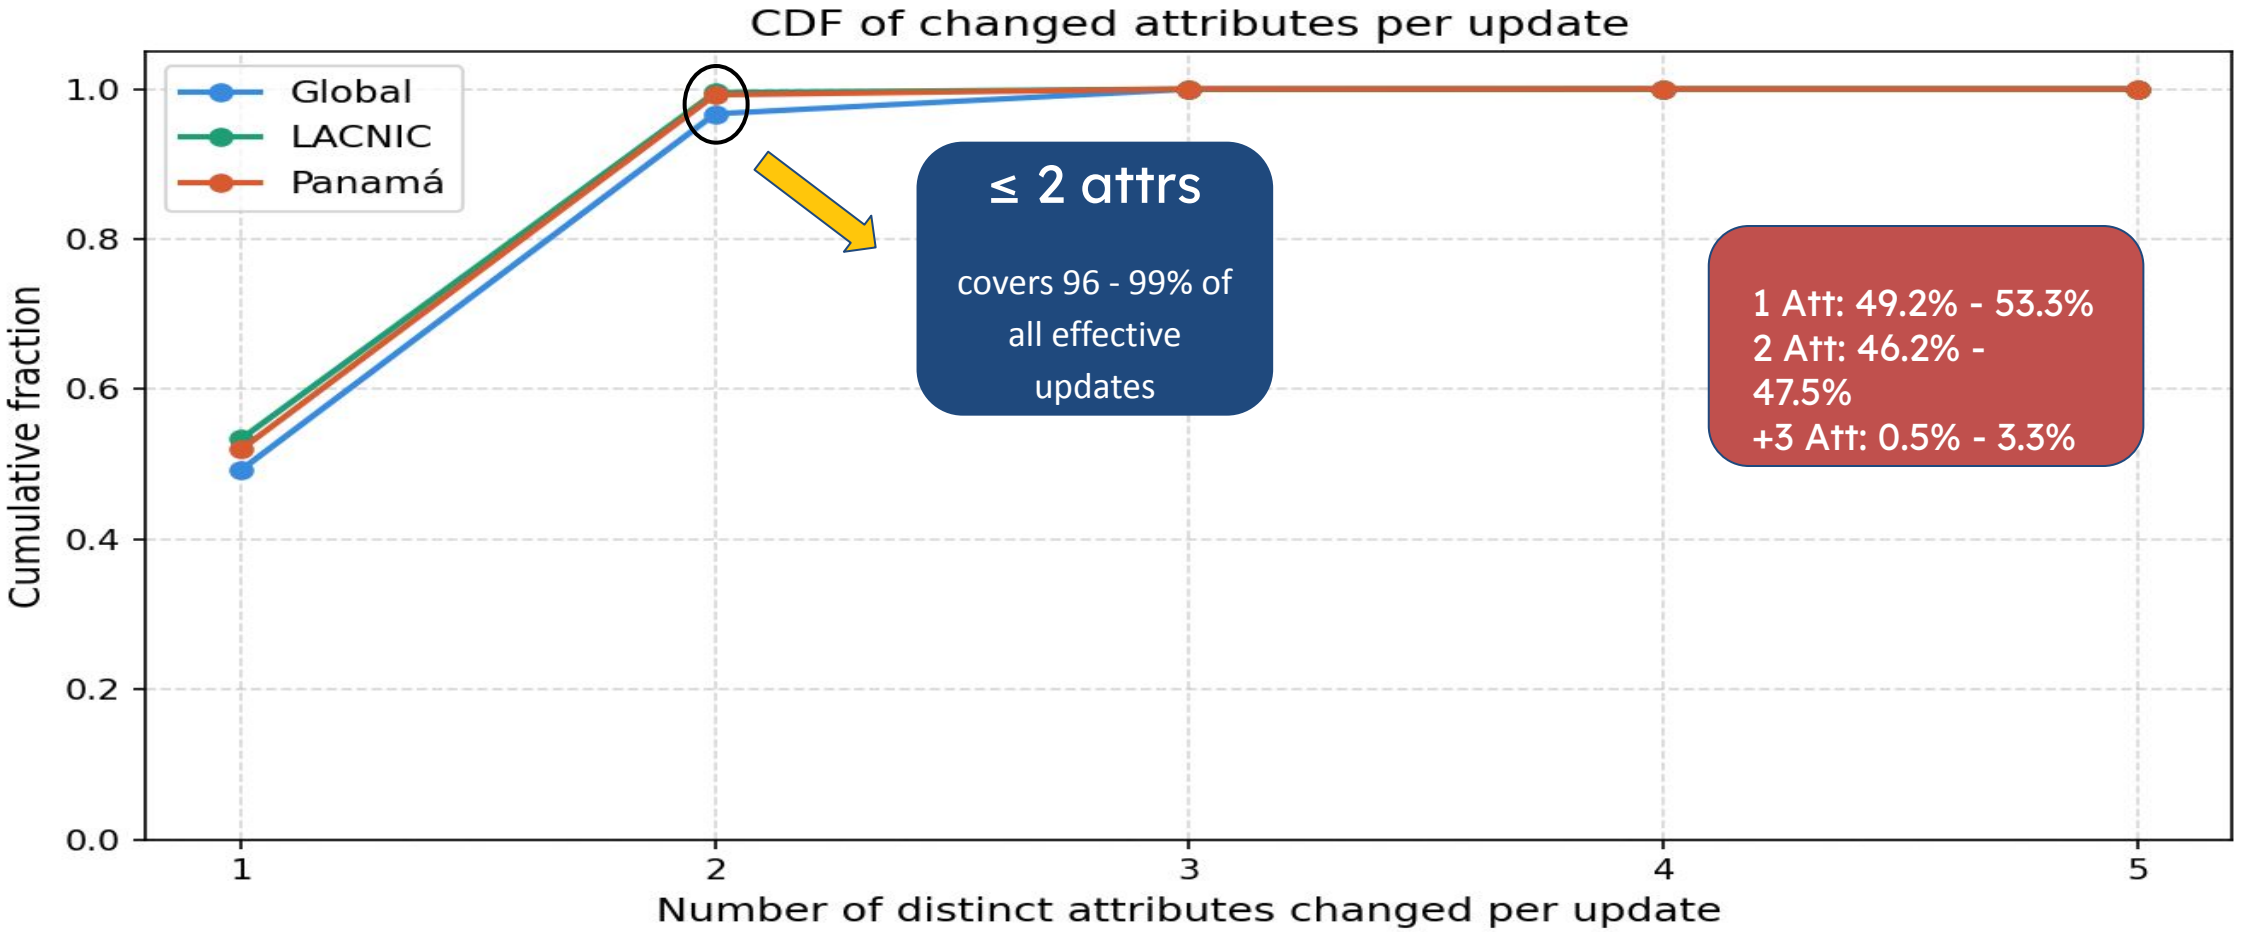

Now, which attributes are behind these changes?

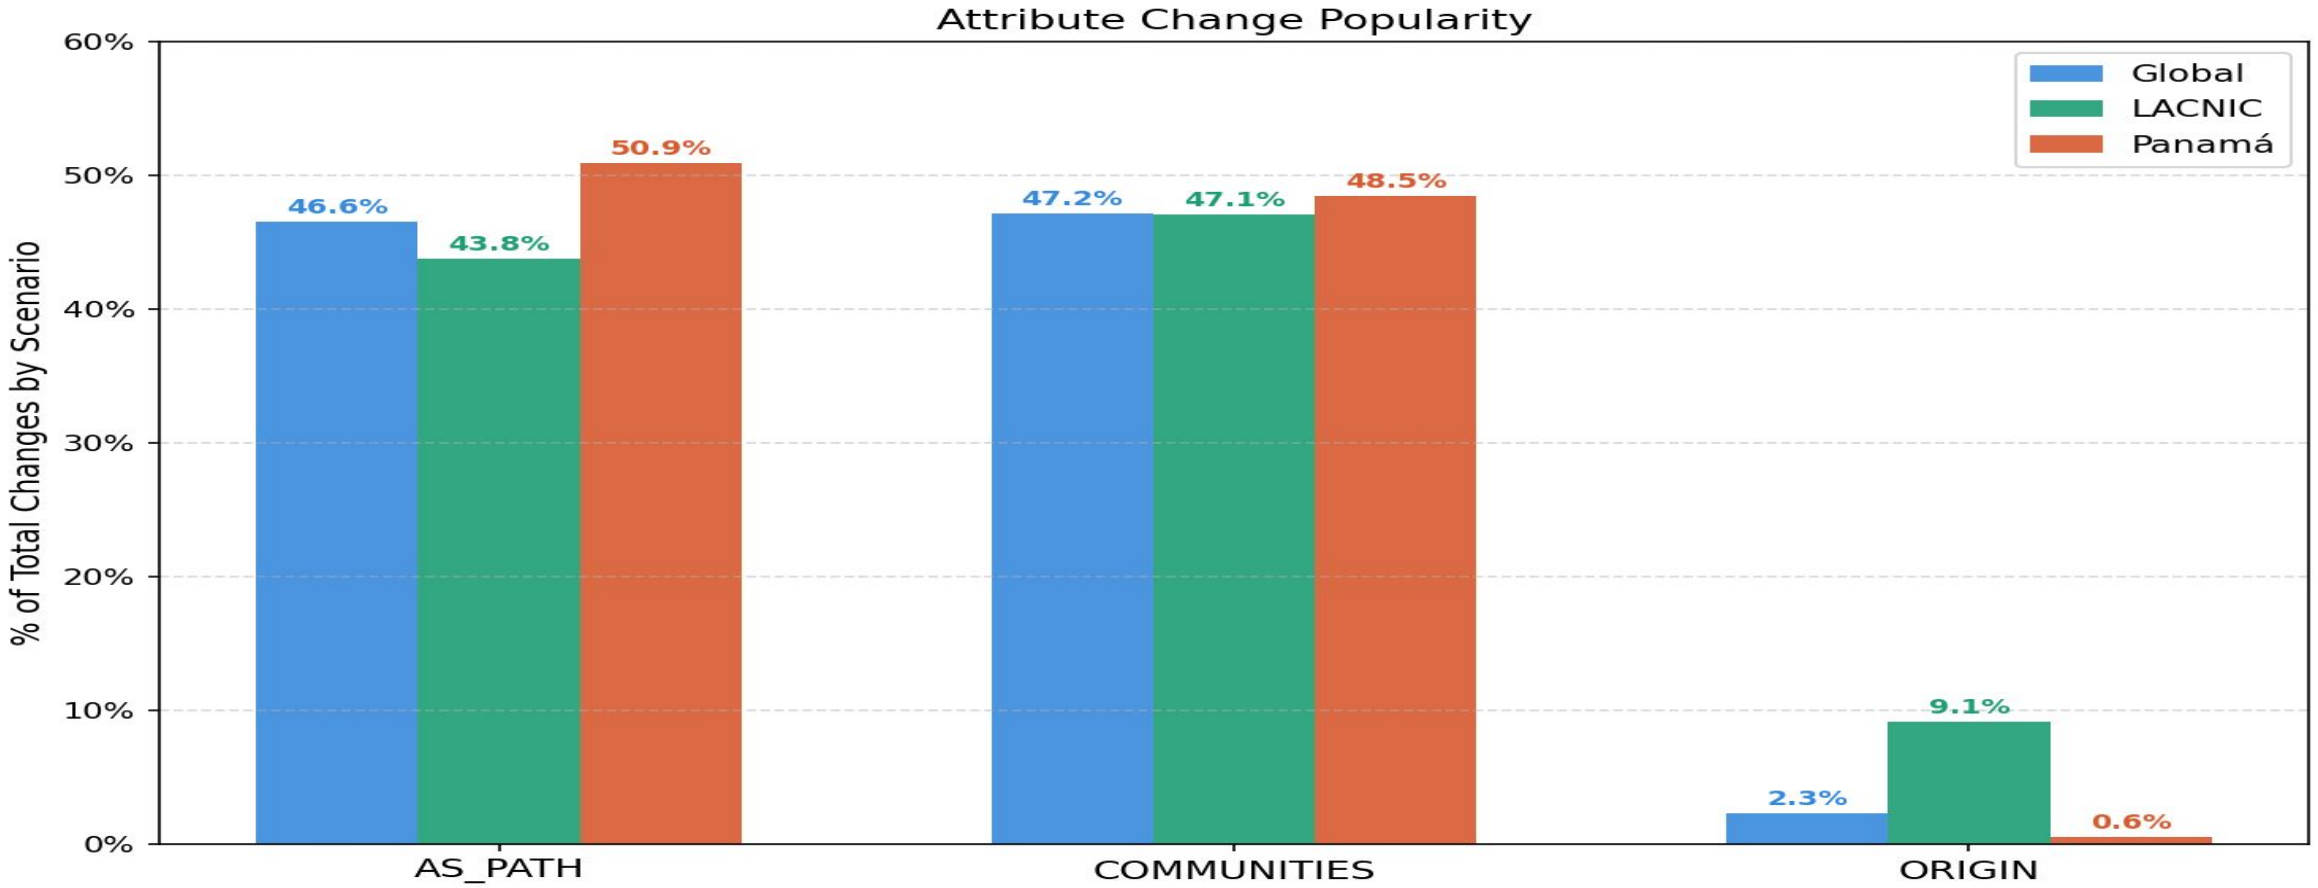

Now, which attributes are behind these changes?

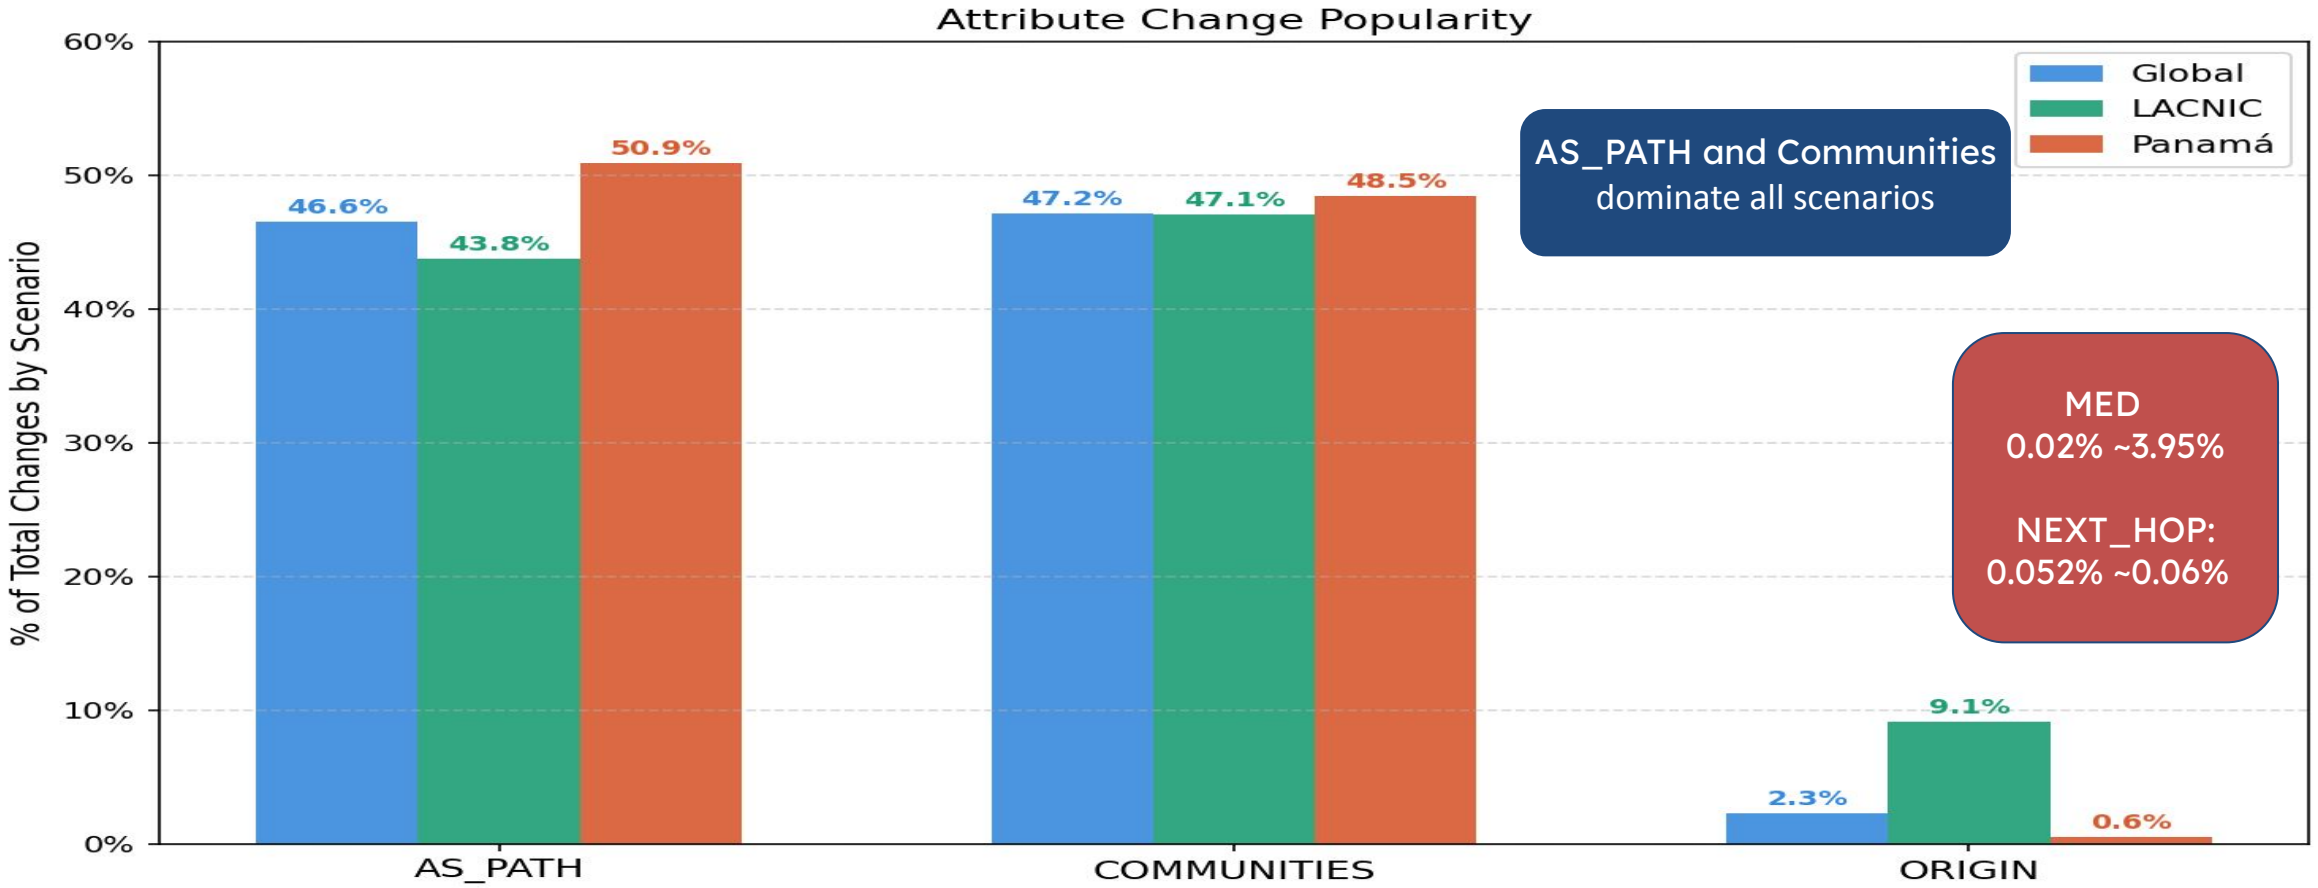

How do we measure propagation?

How do we measure propagation?

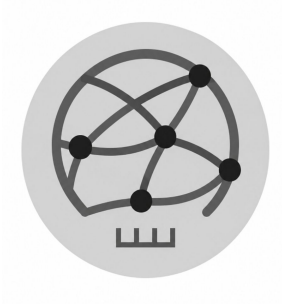

A prefix receives an effective update

How do we measure propagation?

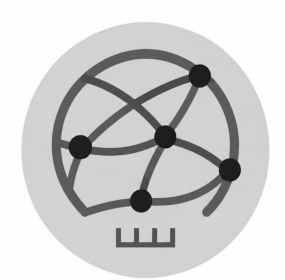

A prefix receives an effective update

We identify all monitors announcing that prefix

How do we measure propagation?

uped within a 120s time windo
per prefix

A prefix receives an effective update

We identify all monitors announcing that prefix

We count how many of them saw the same change

lacnic

How do we measure propagation?

uped within a 120s time windo
per prefix

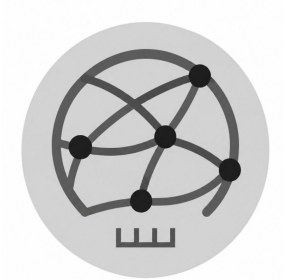

A prefix receives an effective update

We identify all monitors announcing that prefix

We count how many of them saw the same change

Aggregated across all prefixes ->CDF

Repeated for every effective update IPv4 and IPv6 analyzed separately
·across all three scenarios

lacnic

Tracking the reach of routing changes

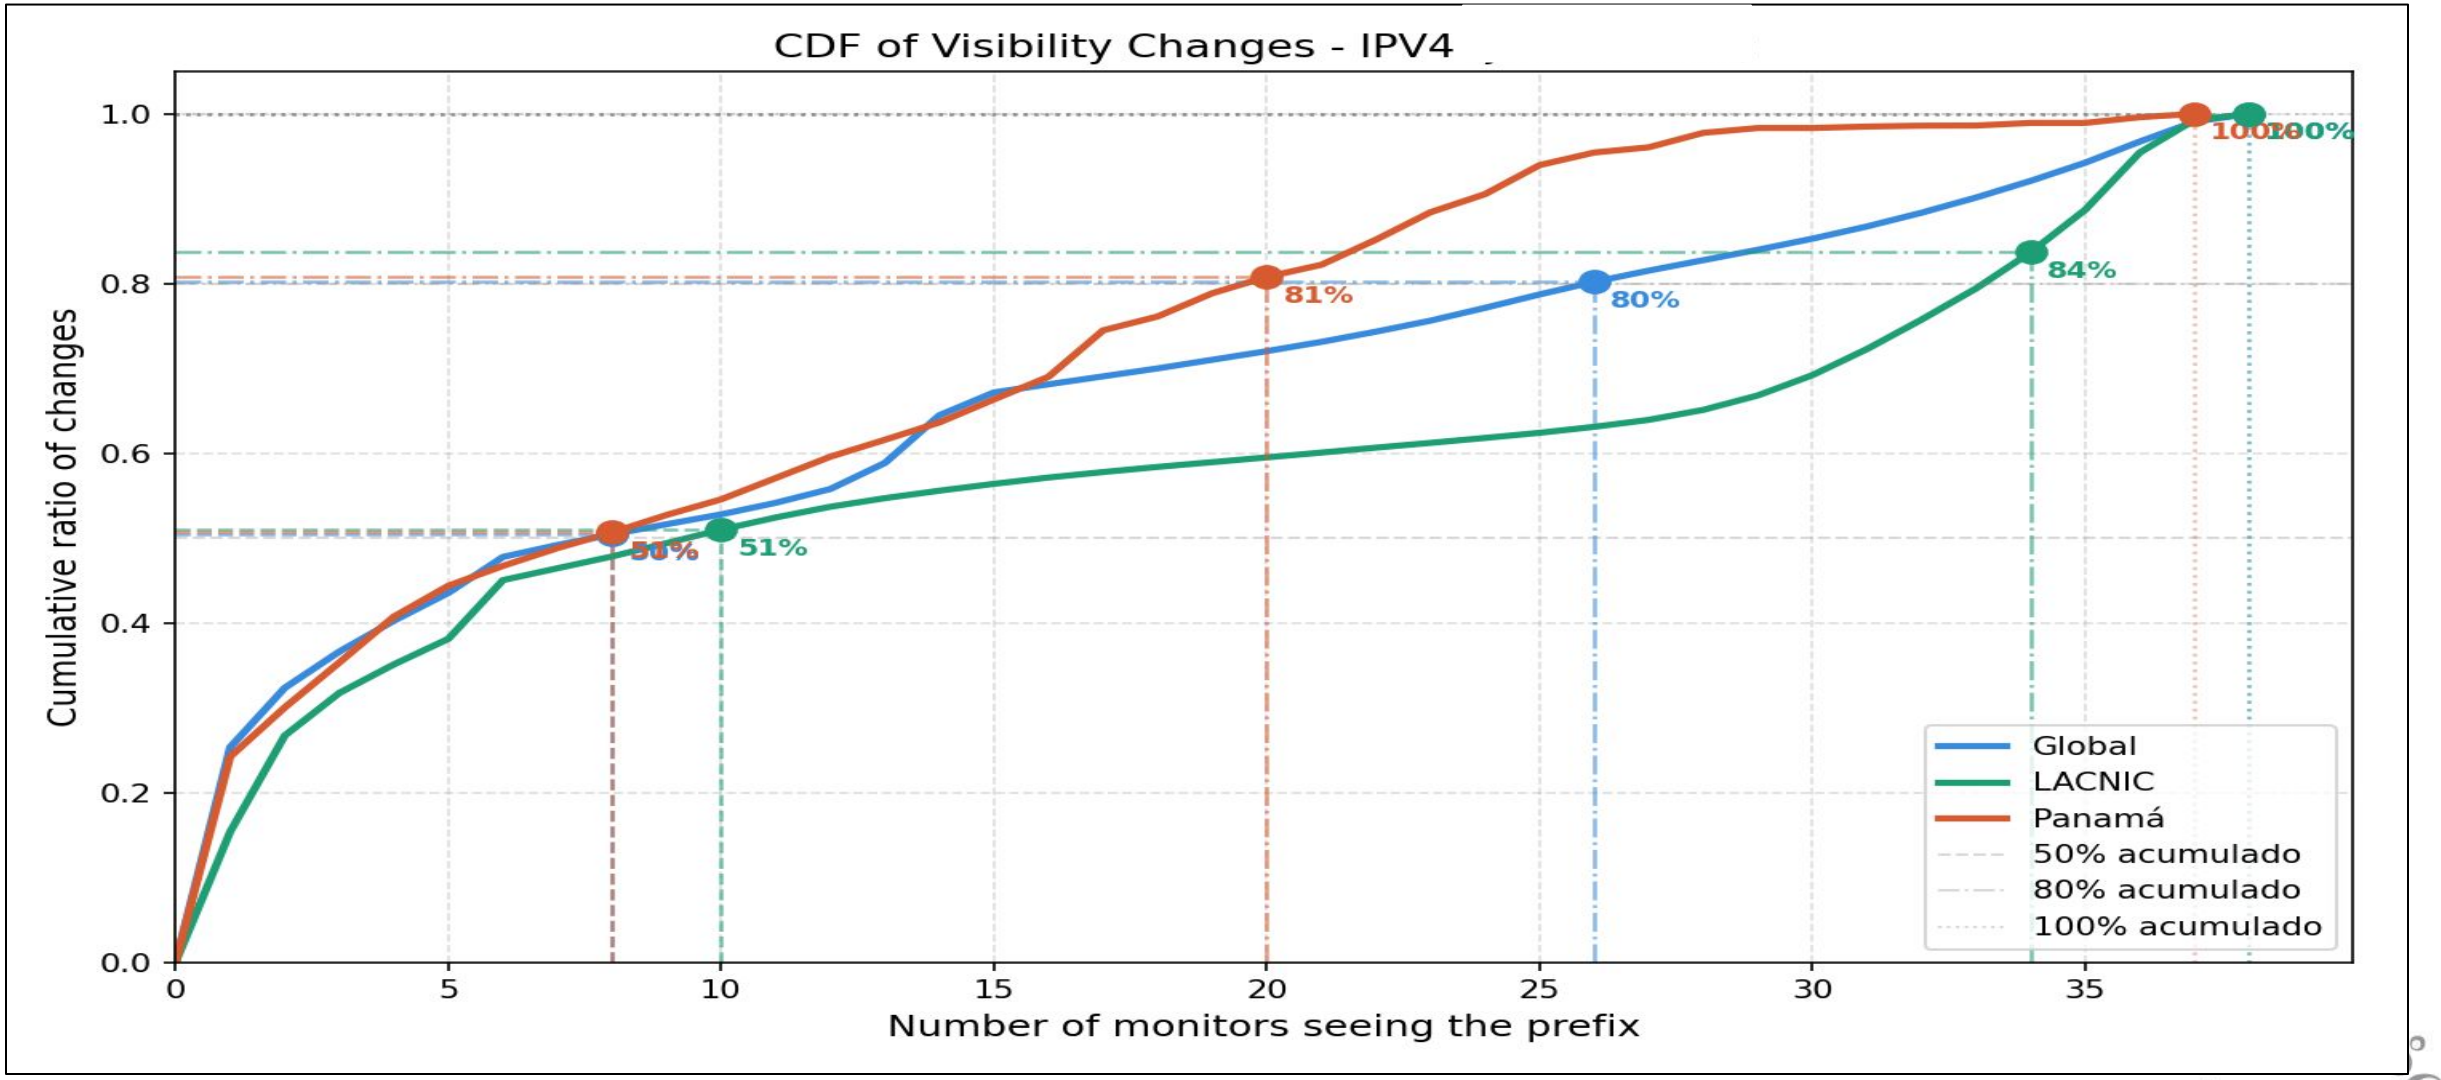

Tracking the reach of routing changes

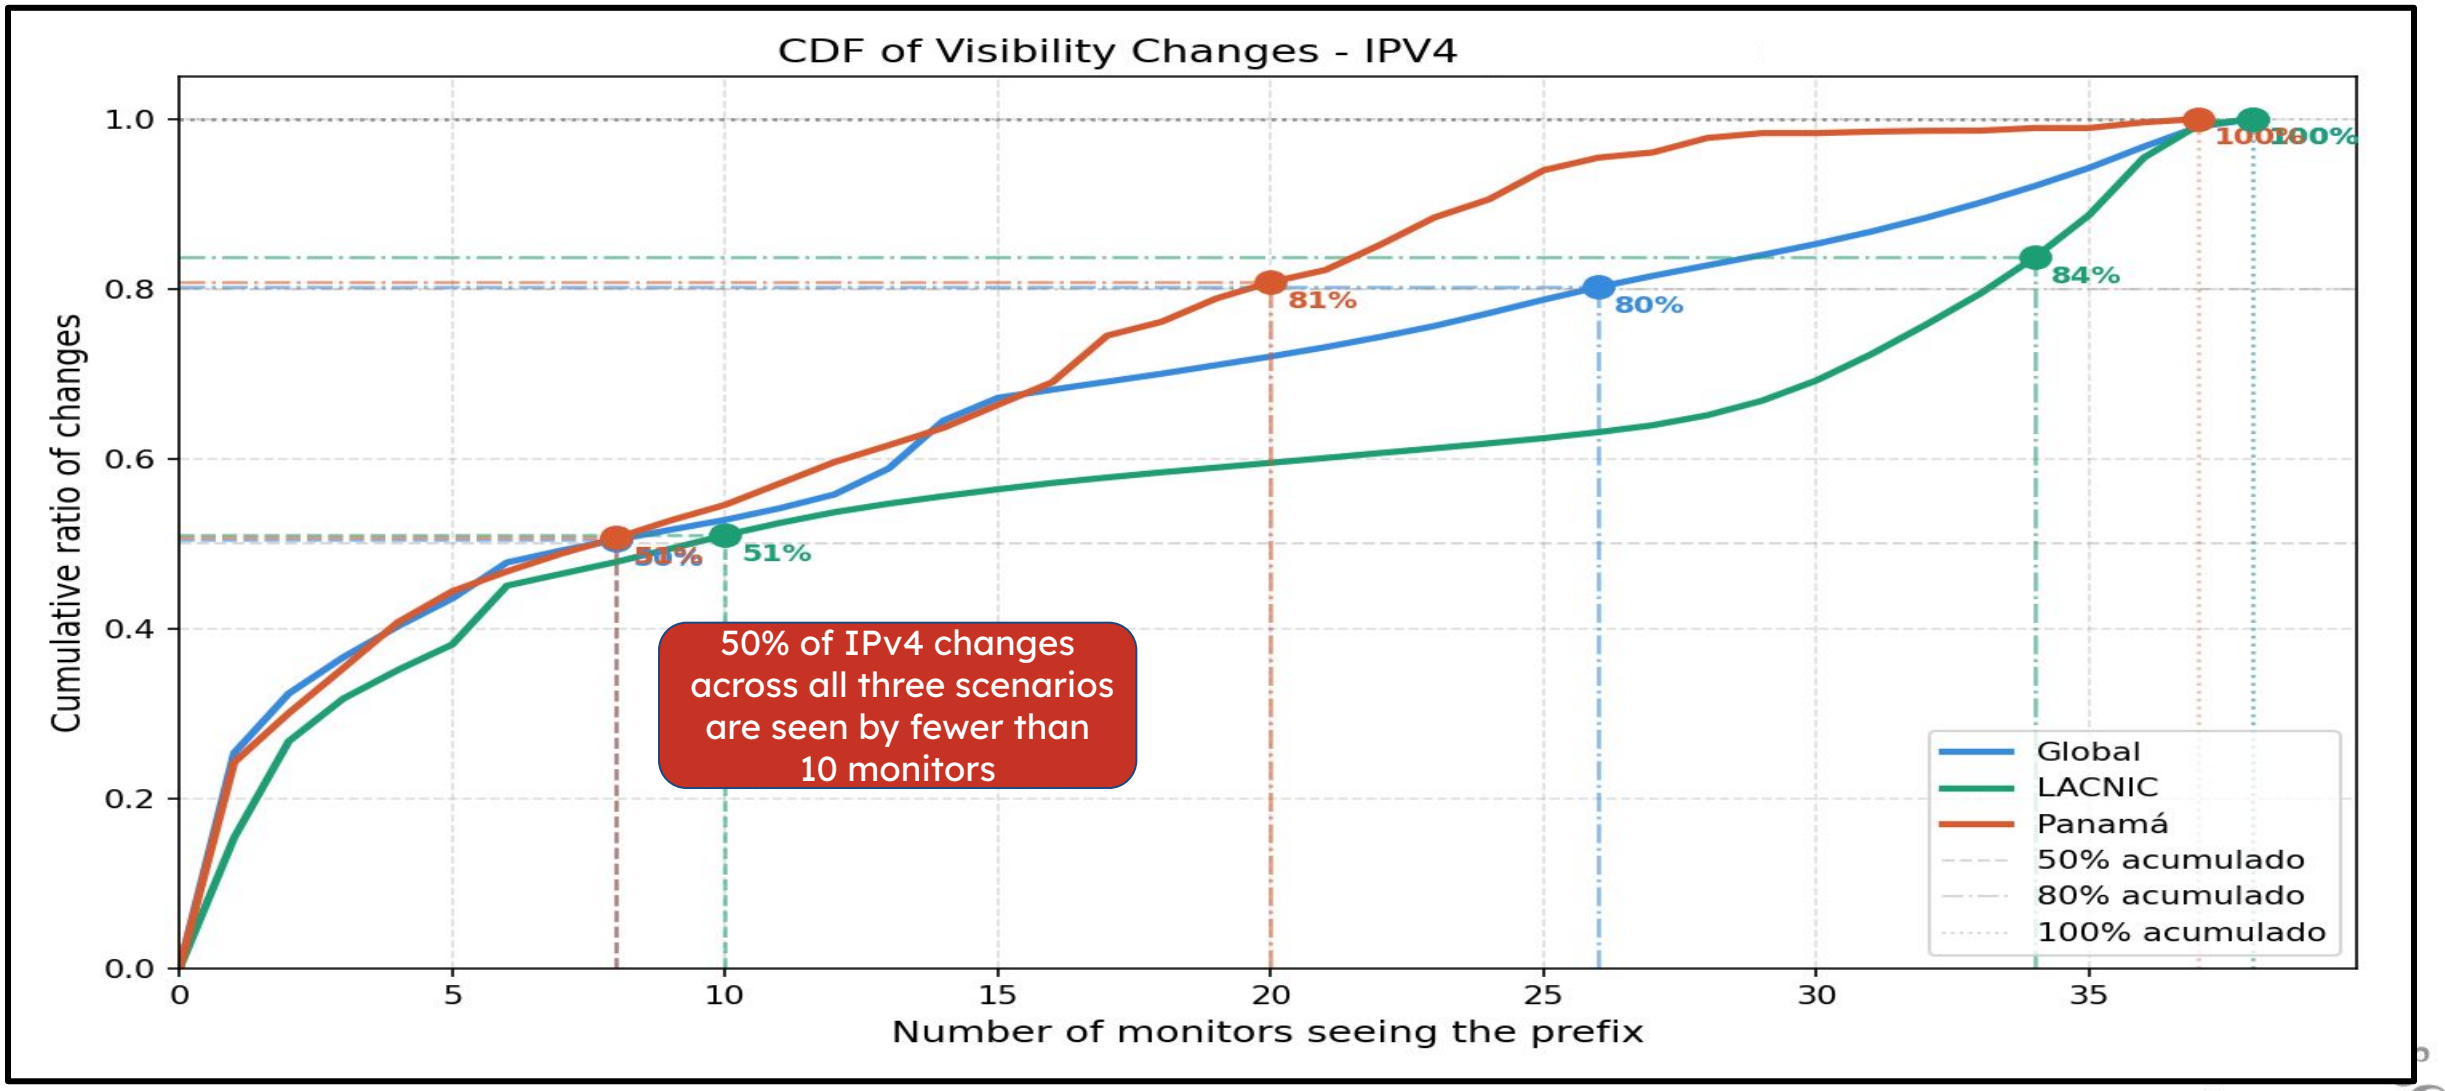

Tracking the reach of routing changes

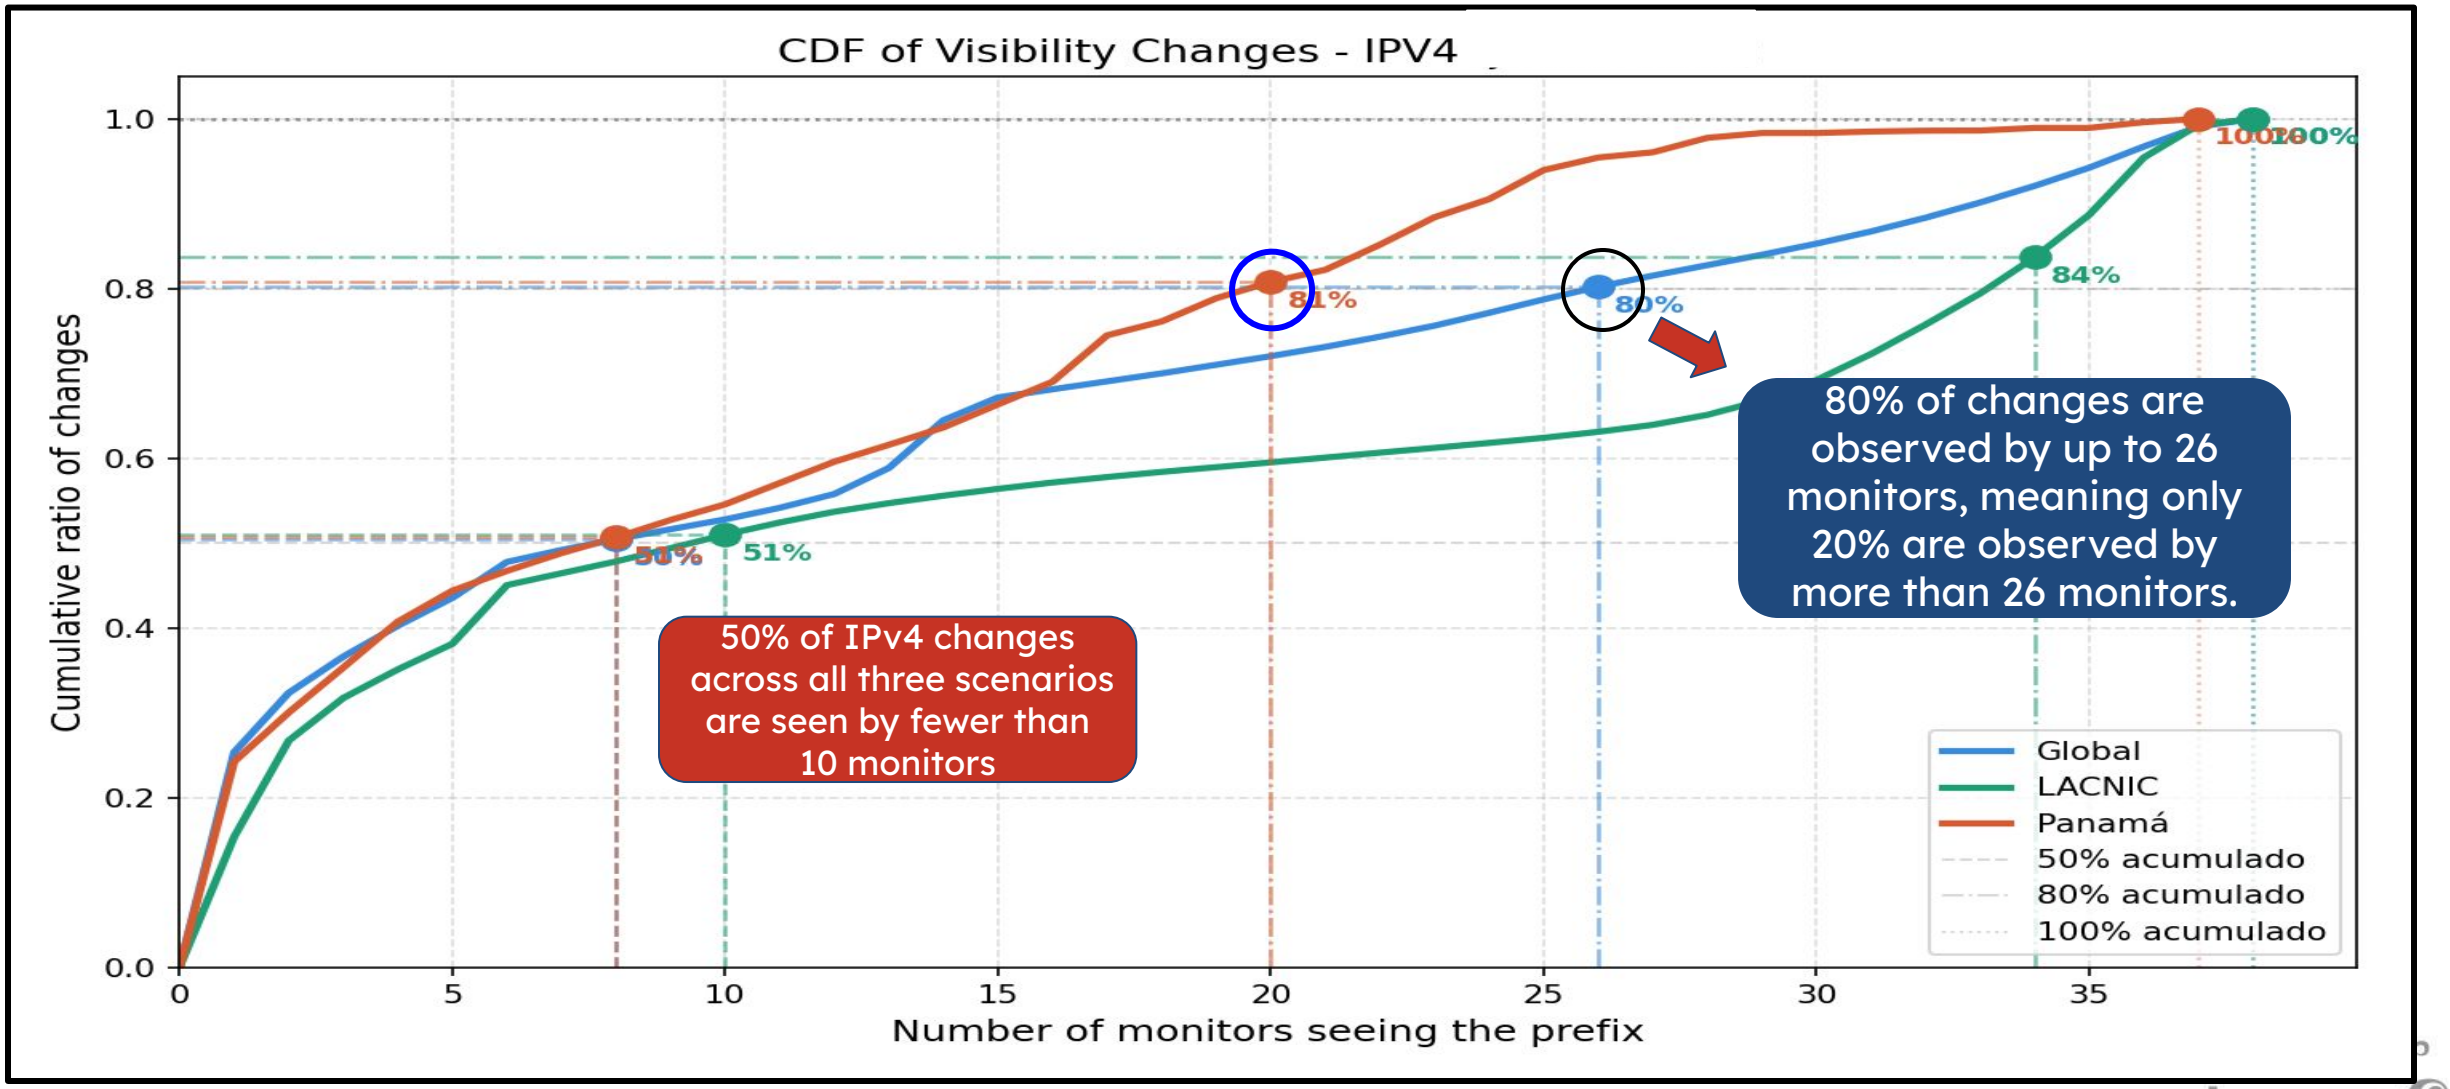

Tracking the reach of routing changes

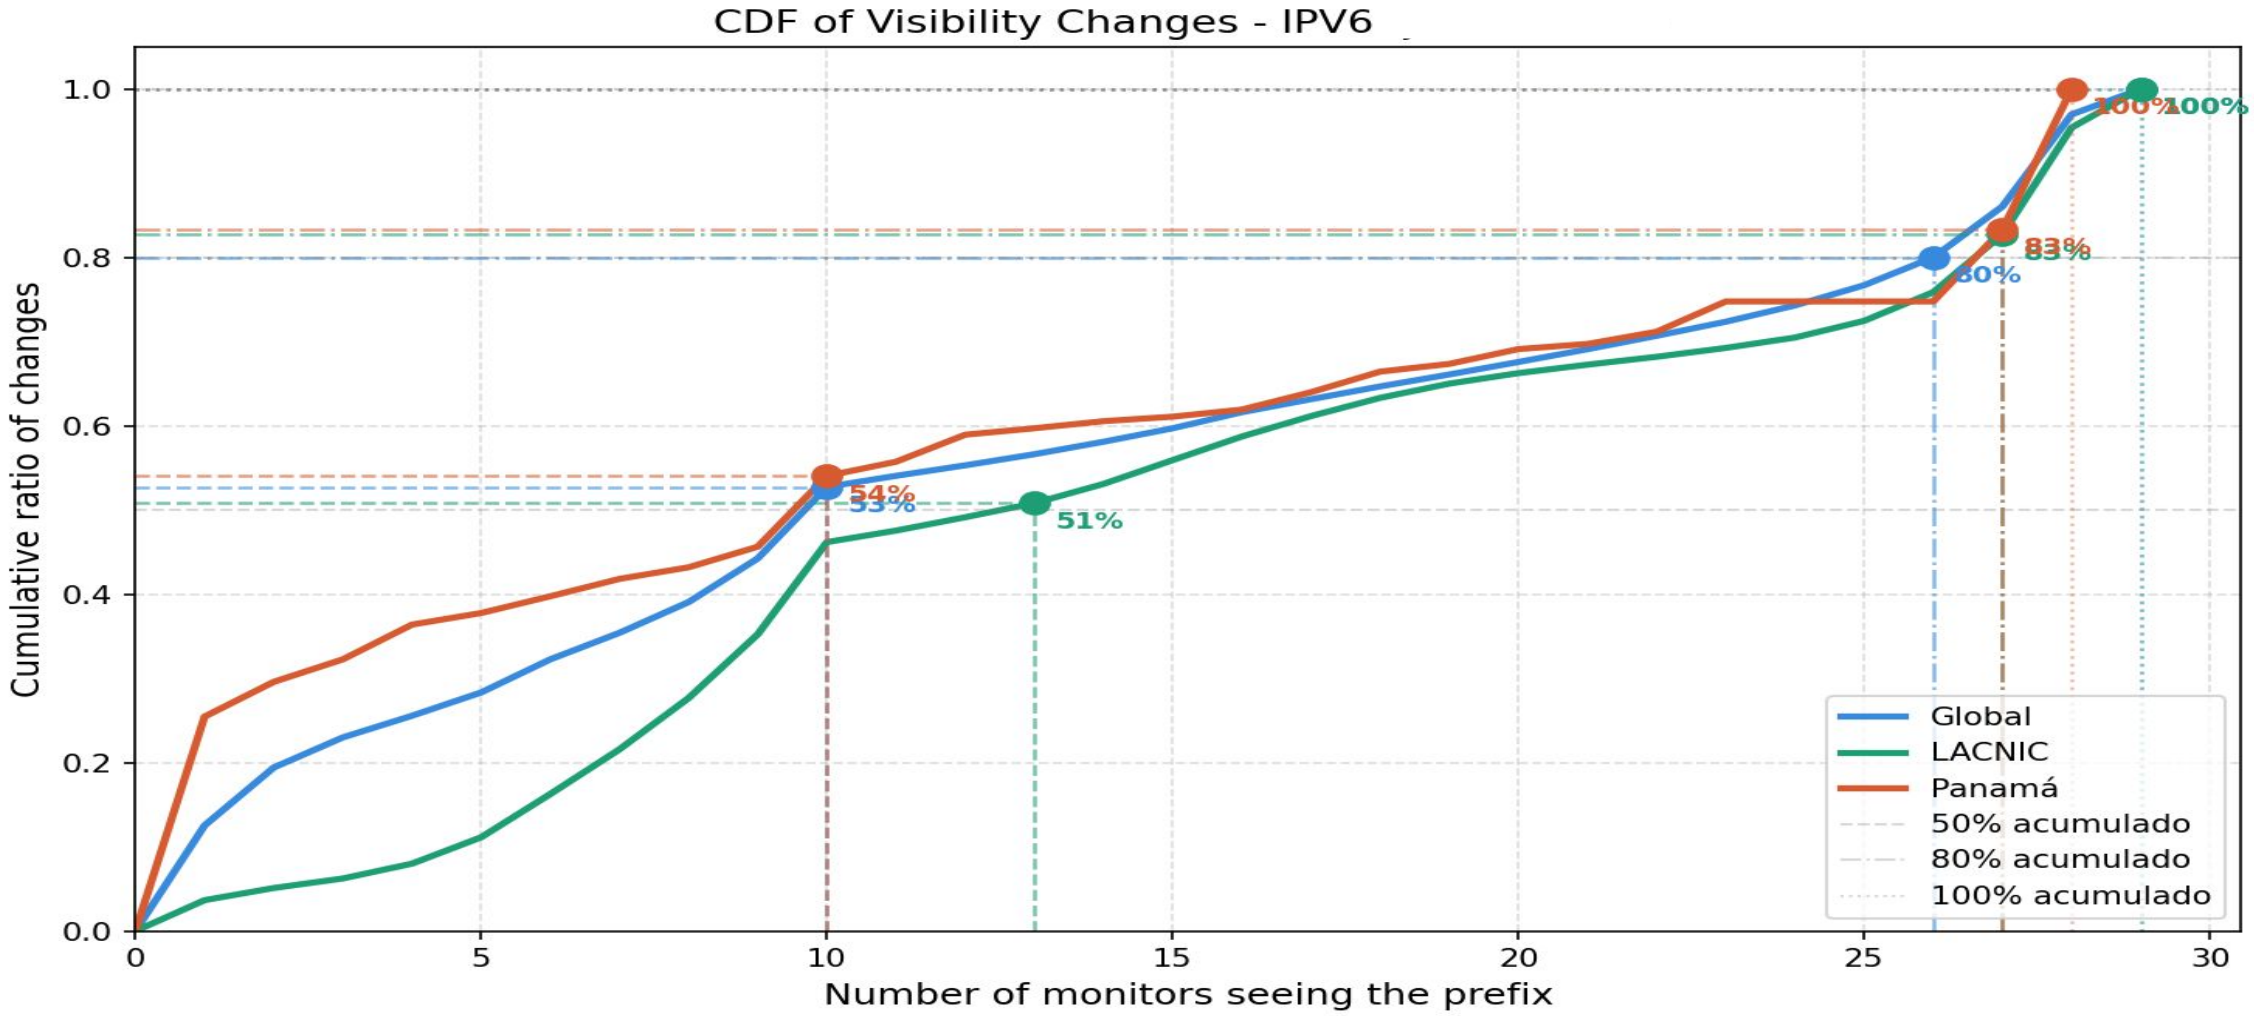

Tracking the reach of routing changes

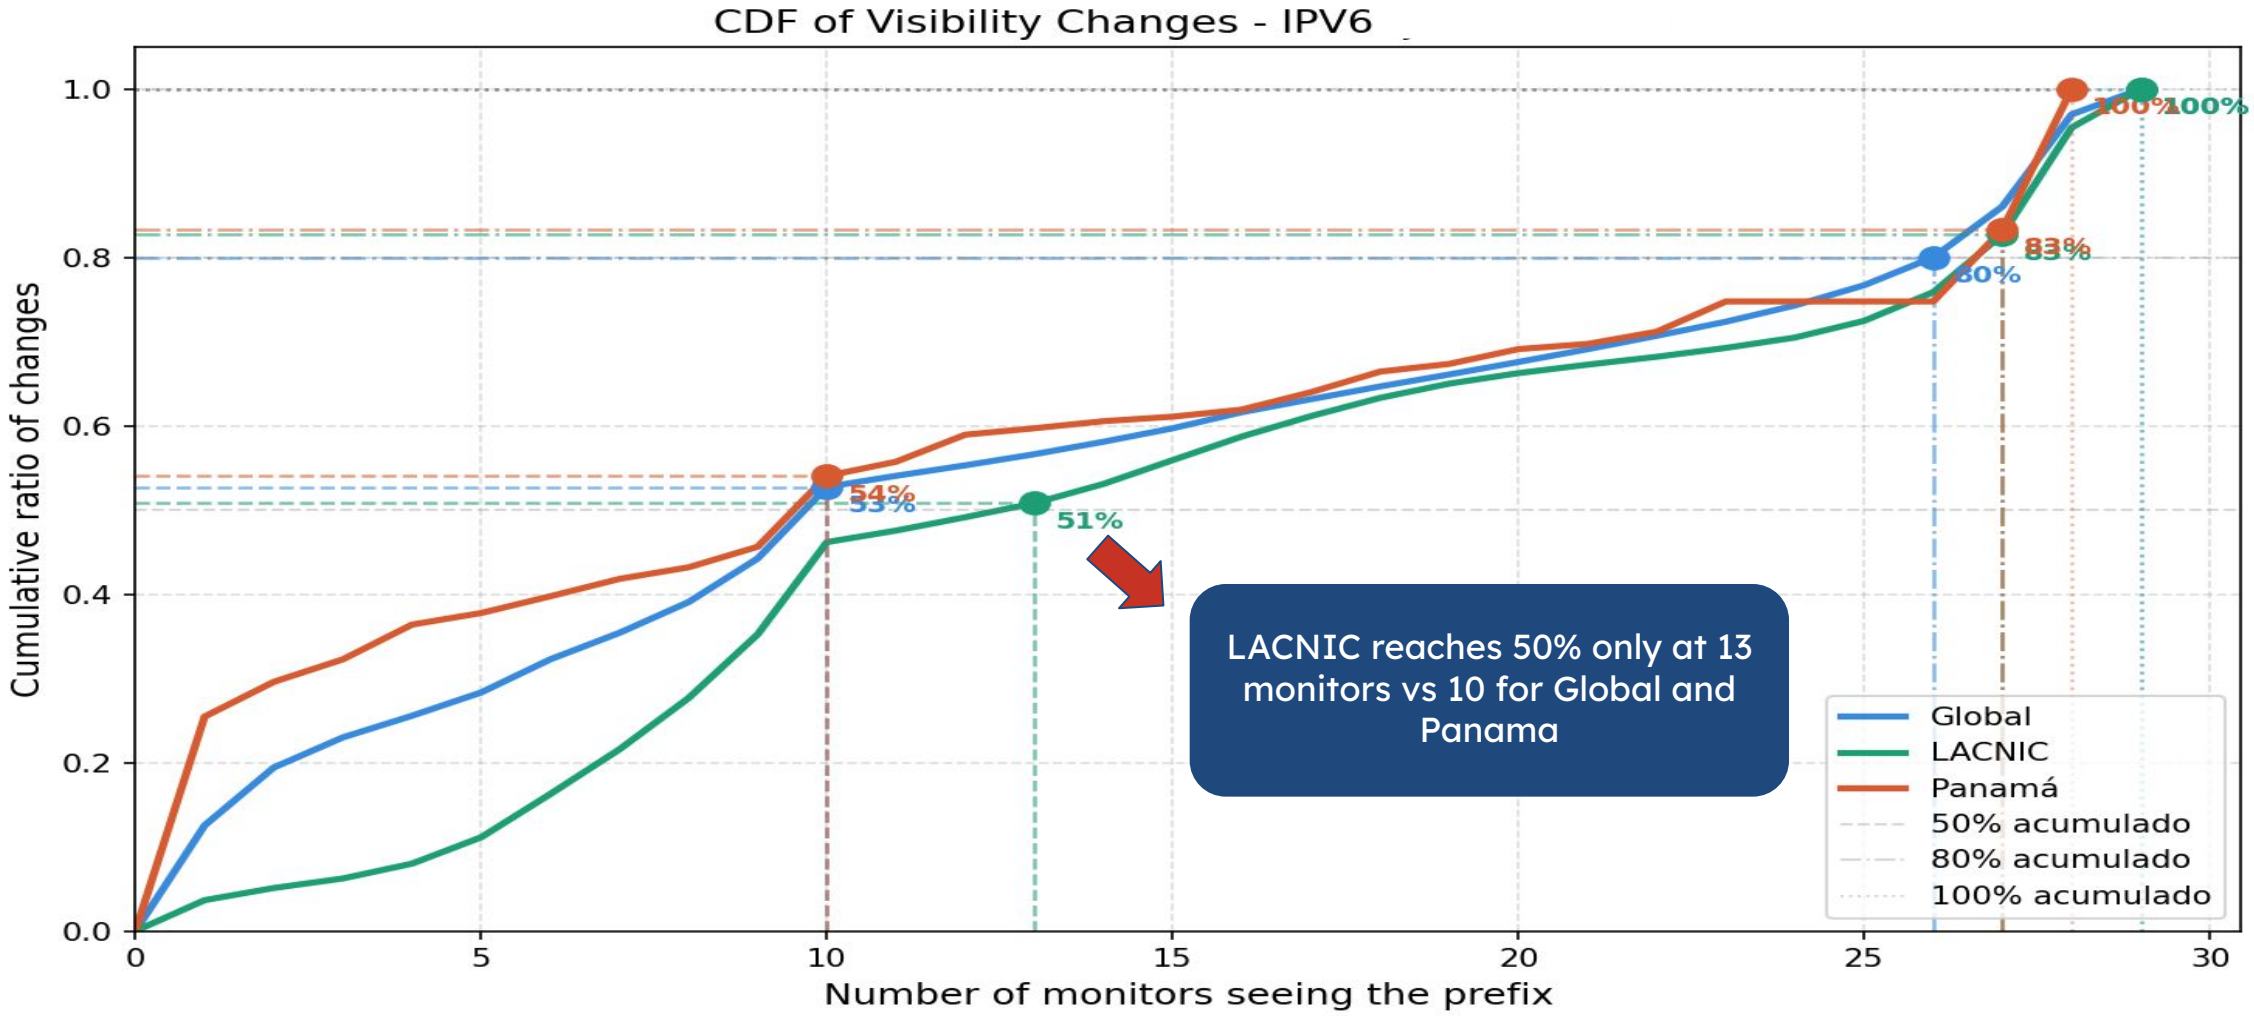

Tracking the reach of routing changes

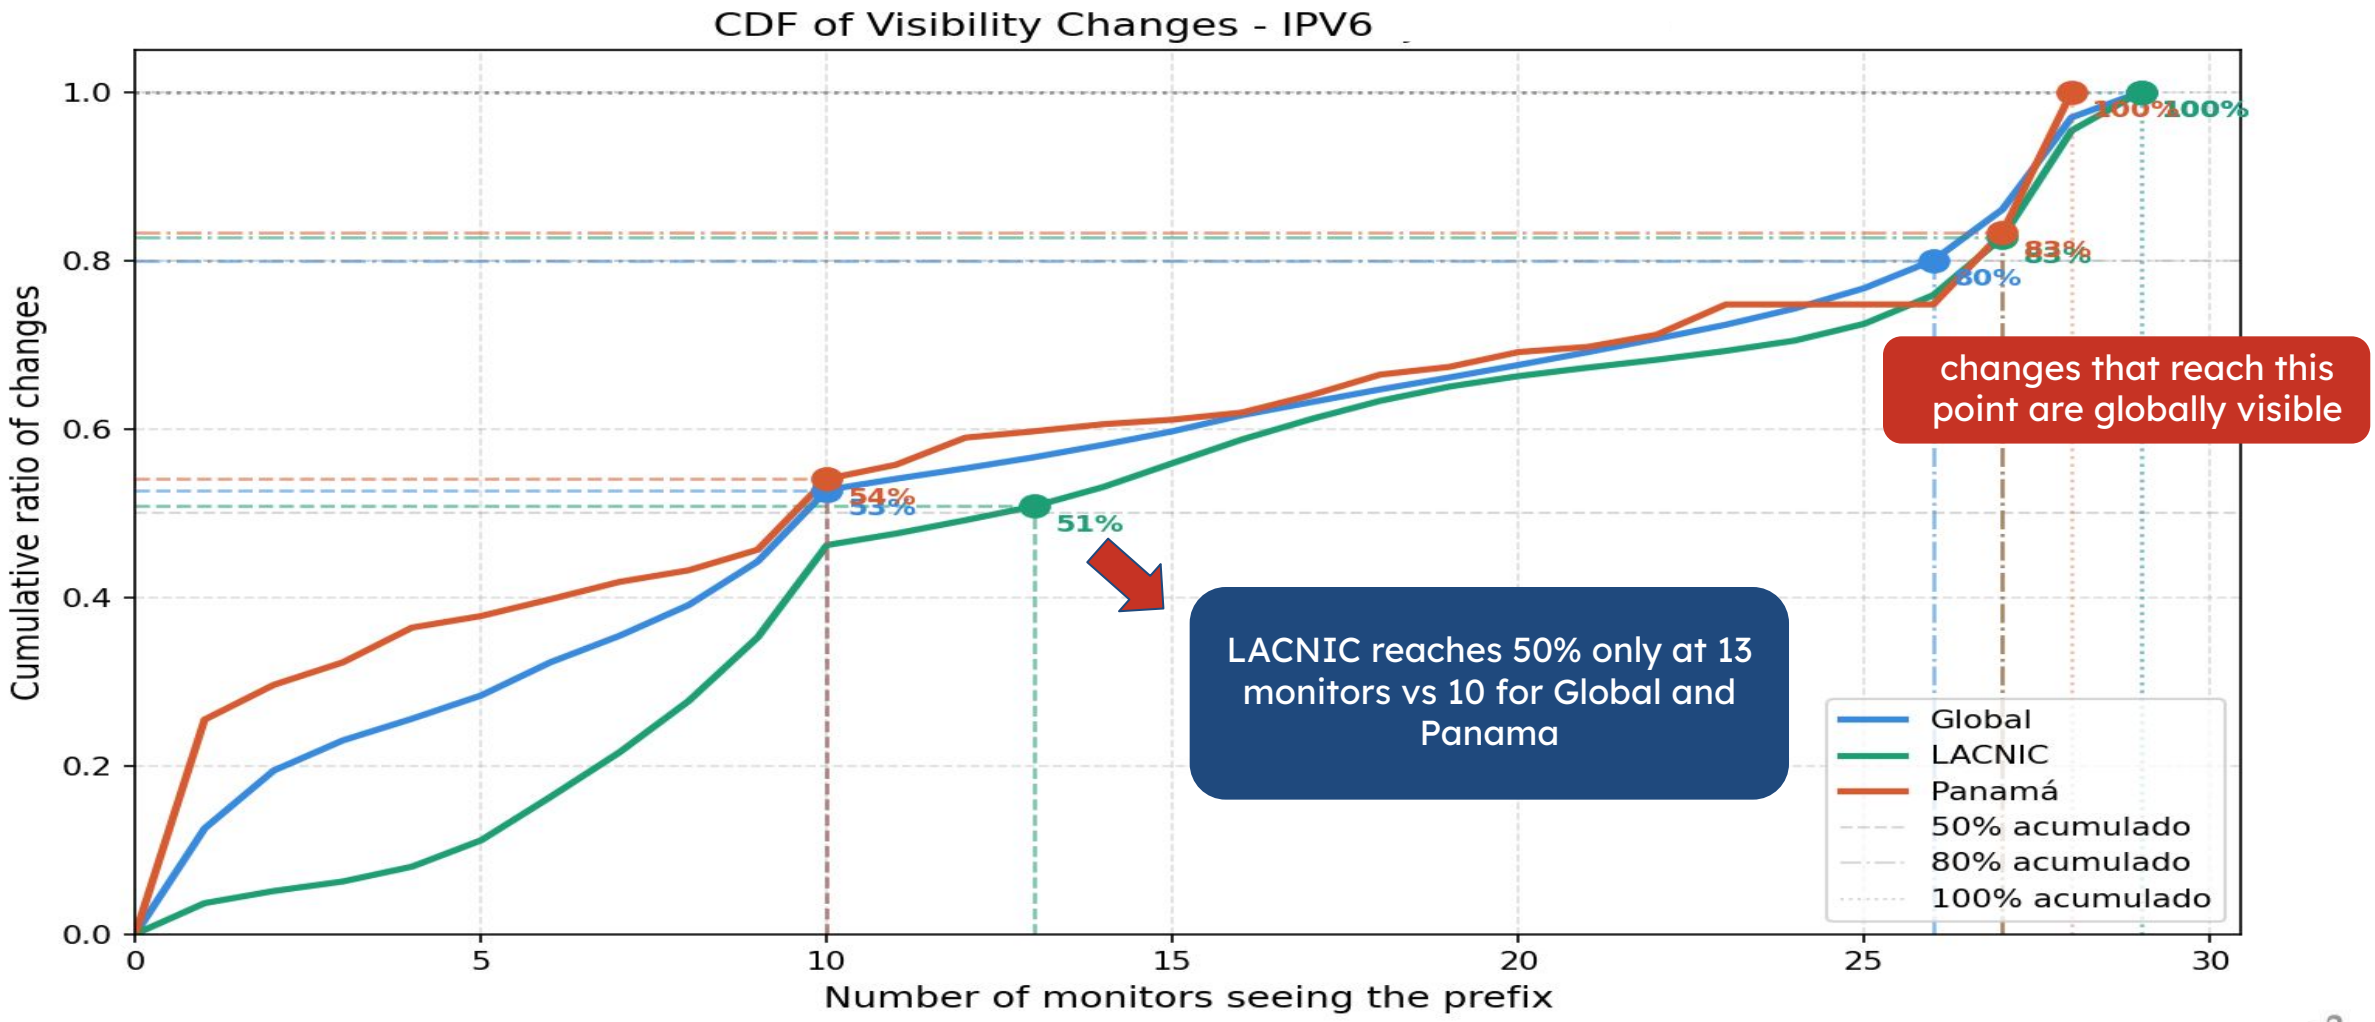

In summary …

In summary …

Q1- Over 60% of BGP updates represent effective routing changes, with LACNIC reaching 81.7%, the highest across all scenarios.

In summary …

Q1- Over 60% of BGP updates represent effective routing changes, with LACNIC reaching 81.7%, the highest across all scenarios.

Q2- AS_PATH and COMMUNITIES dominate attribute changes (~95%), with updates modifying at most 2 attributes in over 96% of cases

In summary …

Q1- Over 60% of BGP updates represent effective routing changes, with LACNIC reaching 81.7%, the highest across all scenarios.

Q2- AS_PATH and COMMUNITIES dominate attribute changes (~95%), with updates modifying at most 2 attributes in over 96% of cases

Q3 - Most BGP routing changes remain locally, 50% of IPv4 changes are seen by ⅓ monitors, while IPv6 changes propagate more widely, reflecting a more interconnected deployment.

In summary …

Q1- Over 60% of BGP updates represent effective routing changes, with LACNIC reaching 81.7%, the highest across all scenarios.

Q2- AS_PATH and COMMUNITIES dominate attribute changes (~95%), with updates modifying at most 2 attributes in over 96% of cases

Q3 - Most BGP routing changes remain locally, 50% of IPv4 changes are seen by ⅓ monitors, while IPv6 changes propagate more widely, reflecting a more interconnected deployment.

What is next …

In summary …

Q1- Over 60% of BGP updates represent effective routing changes, with LACNIC reaching 81.7%, the highest across all scenarios.

Q2- AS_PATH and COMMUNITIES dominate attribute changes (~95%), with updates modifying at most 2 attributes in over 96% of cases

Q3 - Most BGP routing changes remain locally, 50% of IPv4 changes are seen by ⅓ monitors, while IPv6 changes propagate more widely, reflecting a more interconnected deployment.

What is next …

Q4 - Who is causing the instability?

lacnic

In summary …

Q1- Over 60% of BGP updates represent effective routing changes, with LACNIC reaching 81.7%, the highest across all scenarios.

Q2- AS_PATH and COMMUNITIES dominate attribute changes (~95%), with updates modifying at most 2 attributes in over 96% of cases

Q3 - Most BGP routing changes remain locally, 50% of IPv4 changes are seen by ⅓ monitors, while IPv6 changes propagate more widely, reflecting a more interconnected deployment.

Questions ?

Milena Bueno
milenabueno@furg.br

What is next …

Q4 - Who is causing the instability?

lacnic