MEASURING ANYCAST

Harder Than It Looks

Sander Steffann, Jan Žorž

Netnod Meeting 2026

What is Anycast?

- Same IP address announced from multiple locations via BGP
- Routers send traffic to the "nearest" instance
- Used for DNS root servers, CDNs, DDoS scrubbing etc.

- In theory: packets go to the closest node
- In practice: BGP has... opinions

"Nearest" is a word that BGP interprets very creatively.

Our Anycast Network

- 11 regions across 5 continents
- 33 anycast nodes (3 per region)
- 6 anycast prefixes (3 IPv4 + 3 IPv6)
- AS 8038 / AS 203993
- BIRD 2.14 on all nodes

APE • BRA • ITA • JAP • JNB

LAX • LJU • MIA • NYC • SIN • SYD

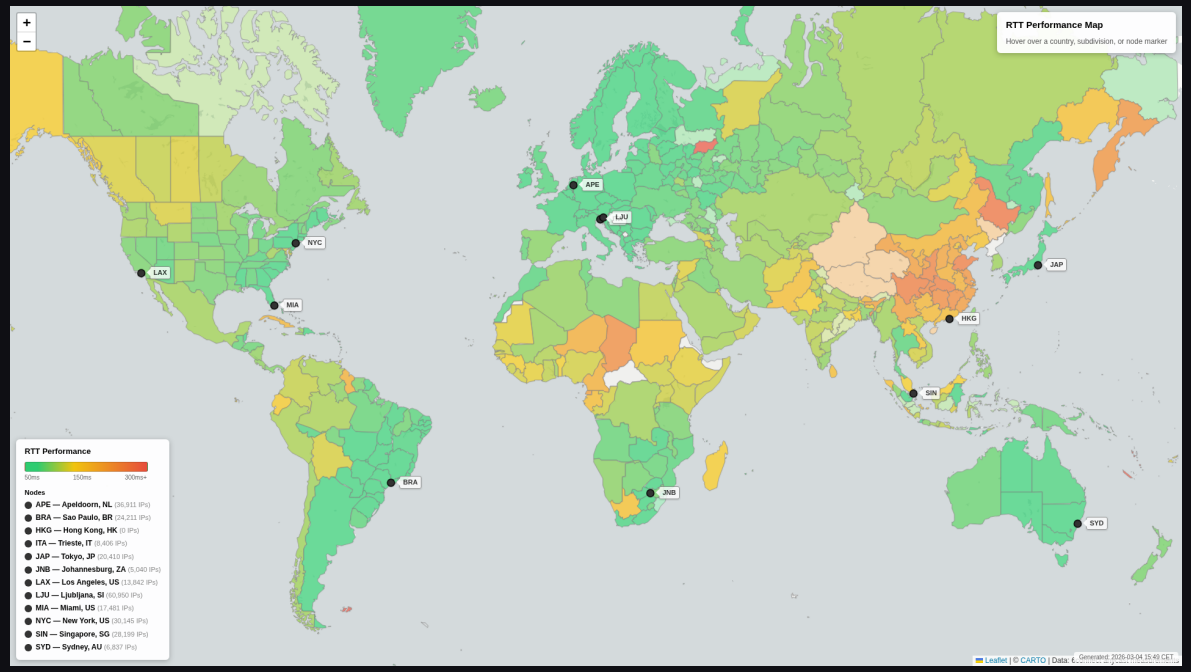

For years, we were flying blind.

We knew the network worked.

DNS queries were answered.

We just couldn't prove it.

Is traffic going to the right node? How fast?

Are packets being dropped? We had no idea.

"How hard can it be to measure anycast?"

— Famous last words

The problem: you're measuring a network that's specifically designed to hide its own topology from you.

The Chicken-and-Egg Problem

- To measure anycast, you send probes FROM your anycast IPs
- Responses come back to... one of your nodes, whichever BGP picks
- You don't control where packets go — that's the whole point
- You need listeners on every node to catch responses

It's turtles all the way down.

First steps

- We don't want to reach others
- We want to know how others reach us!

- How to make others send us traffic?

- Anycast and ICMP are usually a difficult combination
- ICMP errors end up at the wrong node…
- But we can use this to our advantage!

Sometimes it's a bug, sometimes it's a feature

The Architecture We Built

Stats nodes (11)

- Send probes using an anycast VIPs as source address
- Run yarrp, fping, dublin-traceroute

NS nodes (33)

- Serve production DNS
- Catch probe responses via BGP routing
- Run muninn listeners/collectors (ICMP + TCP)

Odin (central server)

- Collects all results, runs analysis pipeline

CHAPTER 1

Casting a Wide Net

ICMP scanning — who's out there?

ICMP Scanning with Yarrp

- Stateless topology scanner at high speed
- 1.3M IPv4 + 378K IPv6 targets, shuffled
- 99,000 probes per second across 11 nodes
- Responses route to nearest node via anycast

What we learned:

- Which node catches each response = catchment
- 1.17 million unique router hops discovered

We also learned yarrp has a kernel bug on Linux 6.9+.
So we patched it. Obviously.

Catchment Map — Who Goes Where?

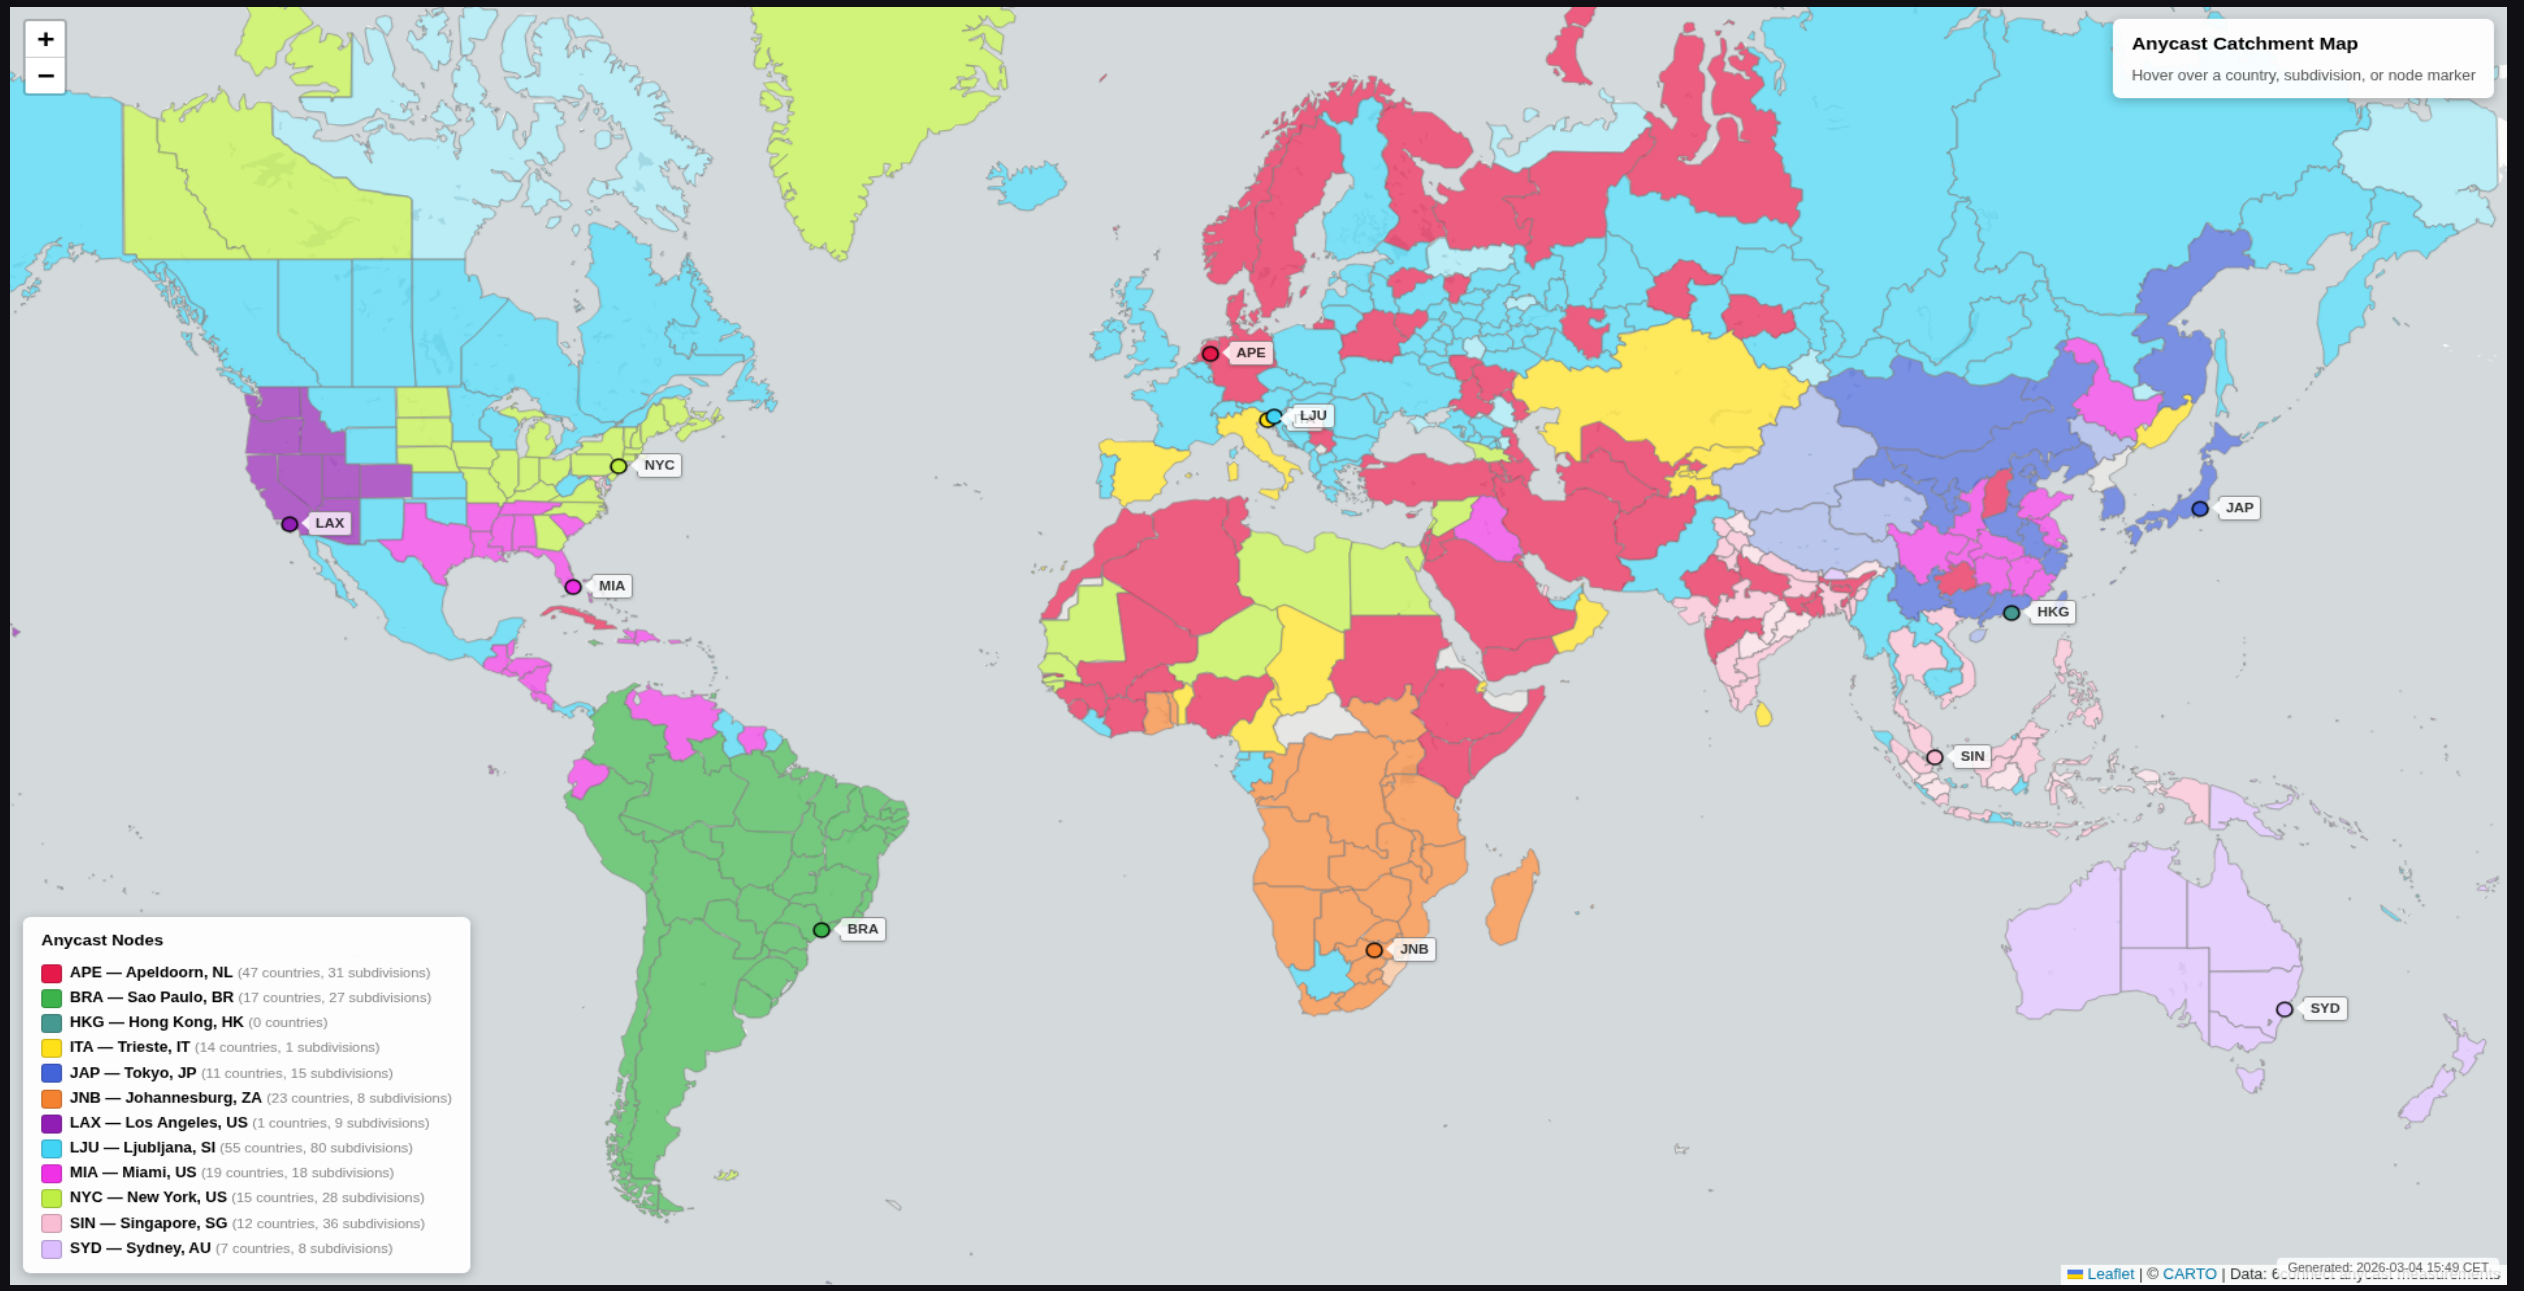

CHAPTER 2

How Fast?

RTT measurements — from 56 minutes of scanning to 5

RTT, Loss & Jitter

- fping in batch mode: 5 probes per target, ~1000 pps
- 775,000 targets across 11 nodes in ~5 minutes
- Per-IP: average RTT, packet loss, jitter

What we learned:

- China routed through Ljubljana (250ms+)
- Some African ISPs: 30%+ packet loss
- South America: surprisingly good (São Paulo)

Old approach: 40 ping subprocesses, 56 minutes.
New approach: one fping process, 5 minutes. Embarrassingly simple.

RTT Performance Map

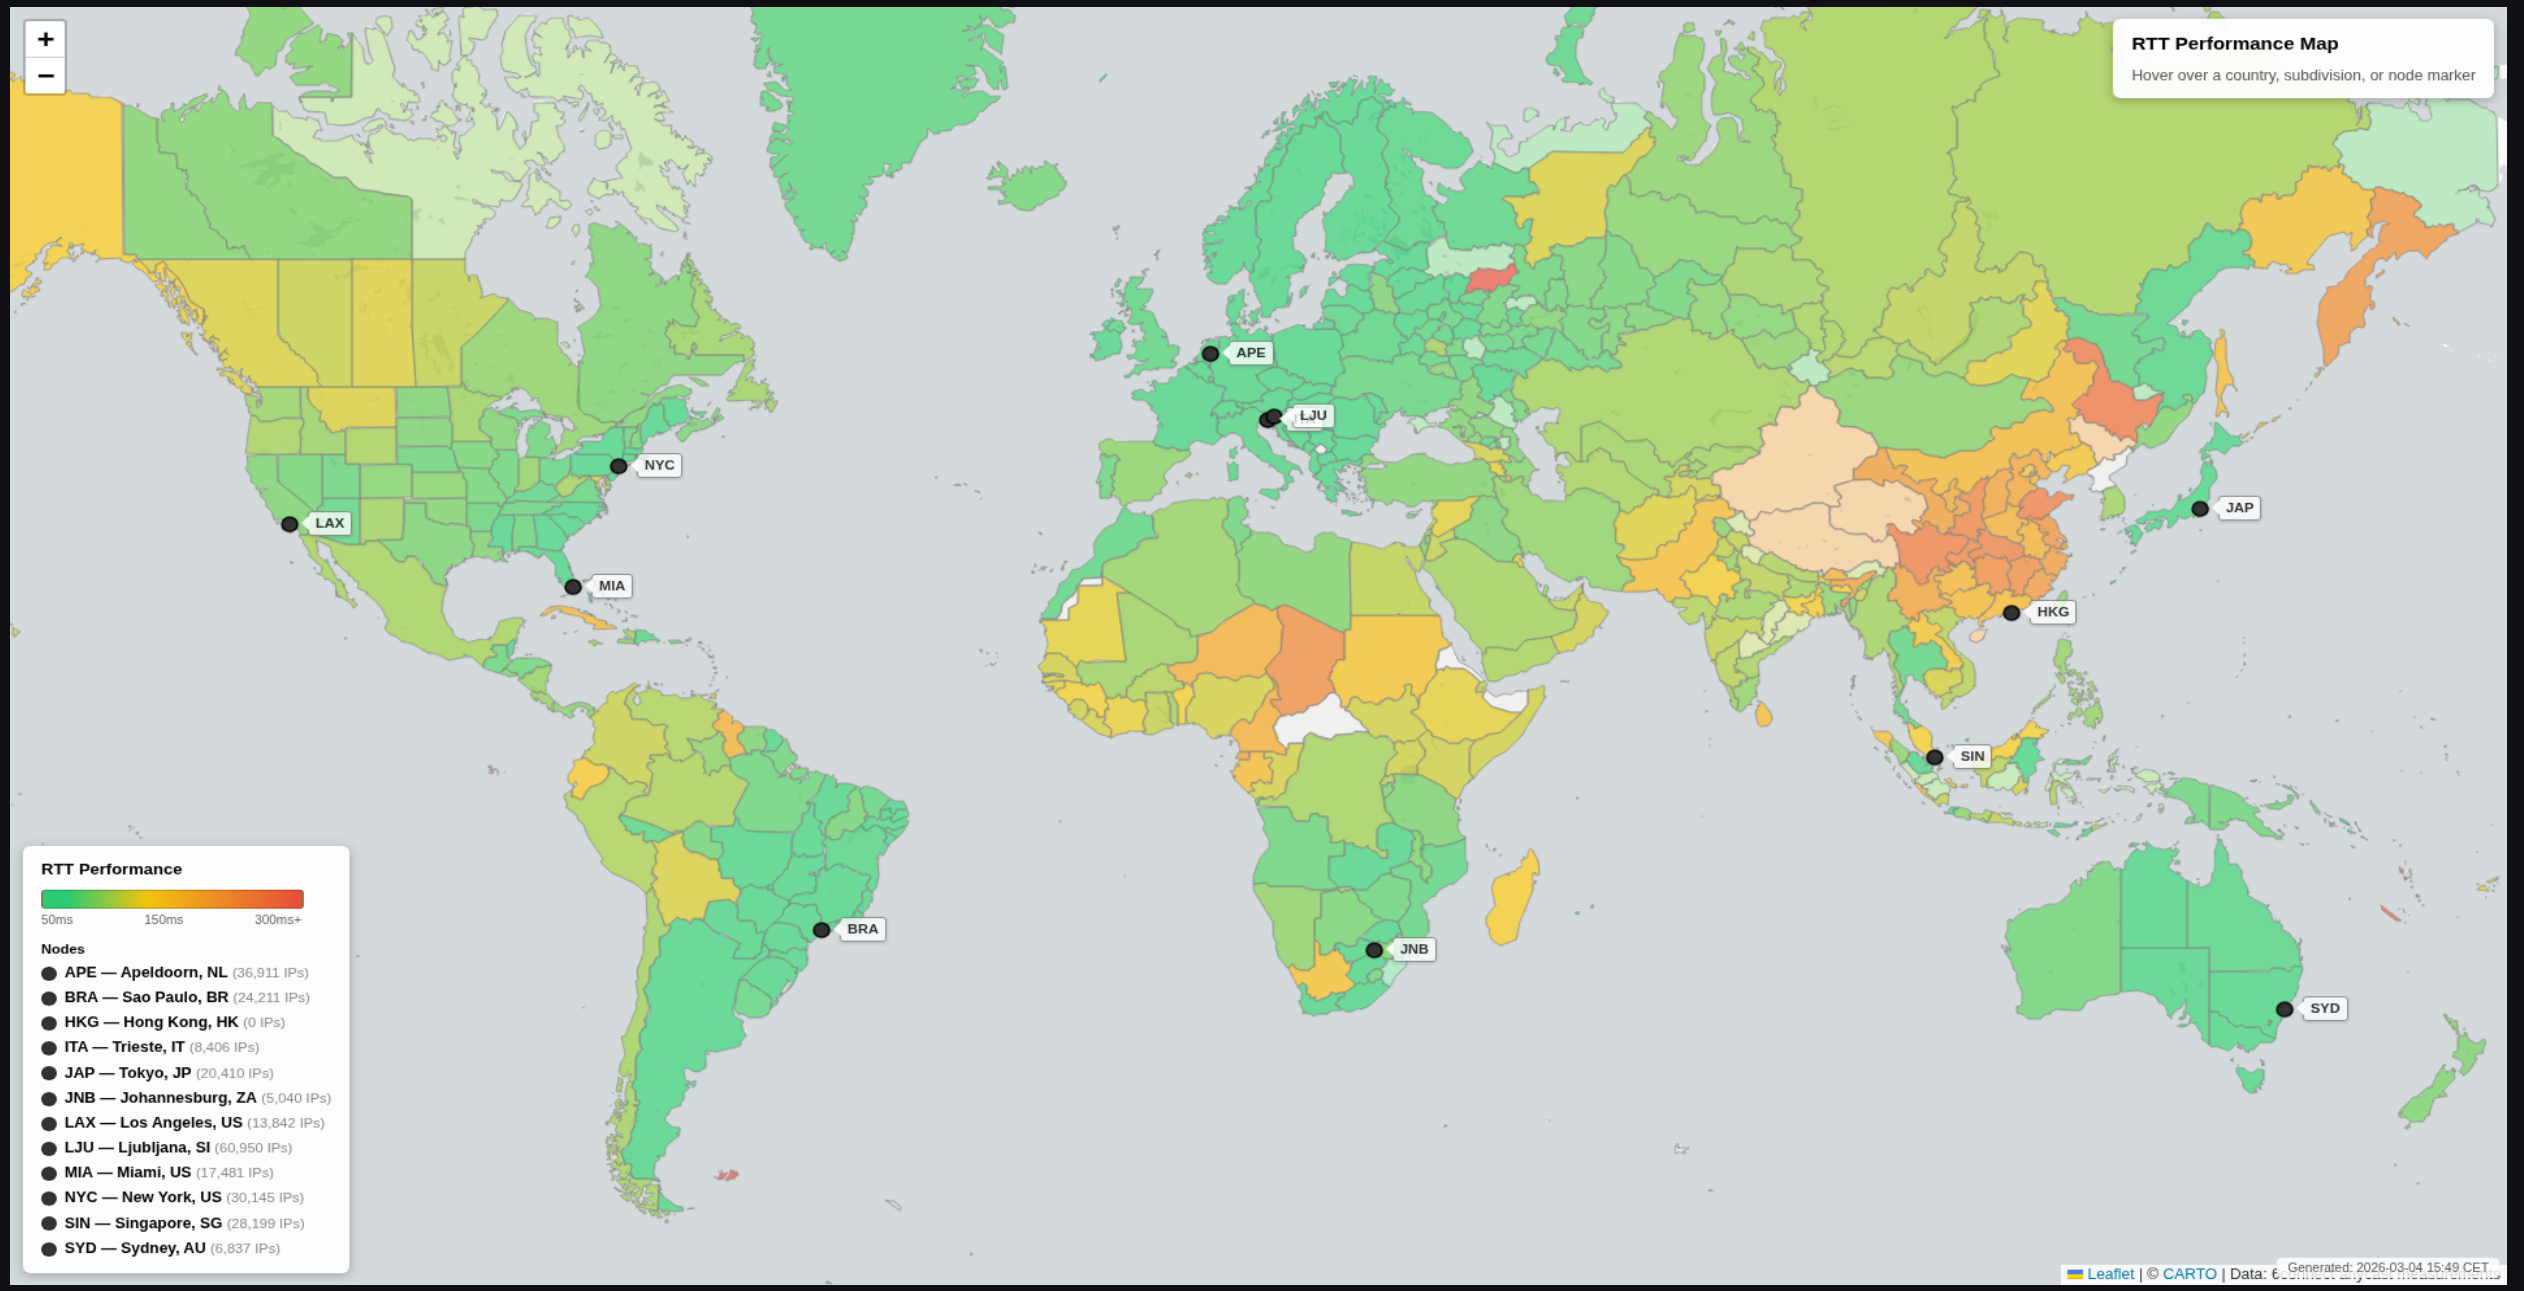

CHAPTER 3

The Silent Killers

PMTU black holes — packets enter, nothing comes out

Path MTU Discovery Black Holes

- 1500-byte packets with Don't Fragment bit
- Routers should reply ICMP "Packet Too Big"
- Some don't. The packet just... disappears.

Three-pass hybrid approach:

- Pass 1: fping 1280B → baseline reachability
- Pass 2: fping 1500B → full-size? (~95% pass)
- Pass 3: ping with DF → ICMP error or silent drop?

fping can't distinguish "got an ICMP error" from "got nothing".
So we use fping for the fast part and real ping for the hard part.

PMTUD Black Hole Map

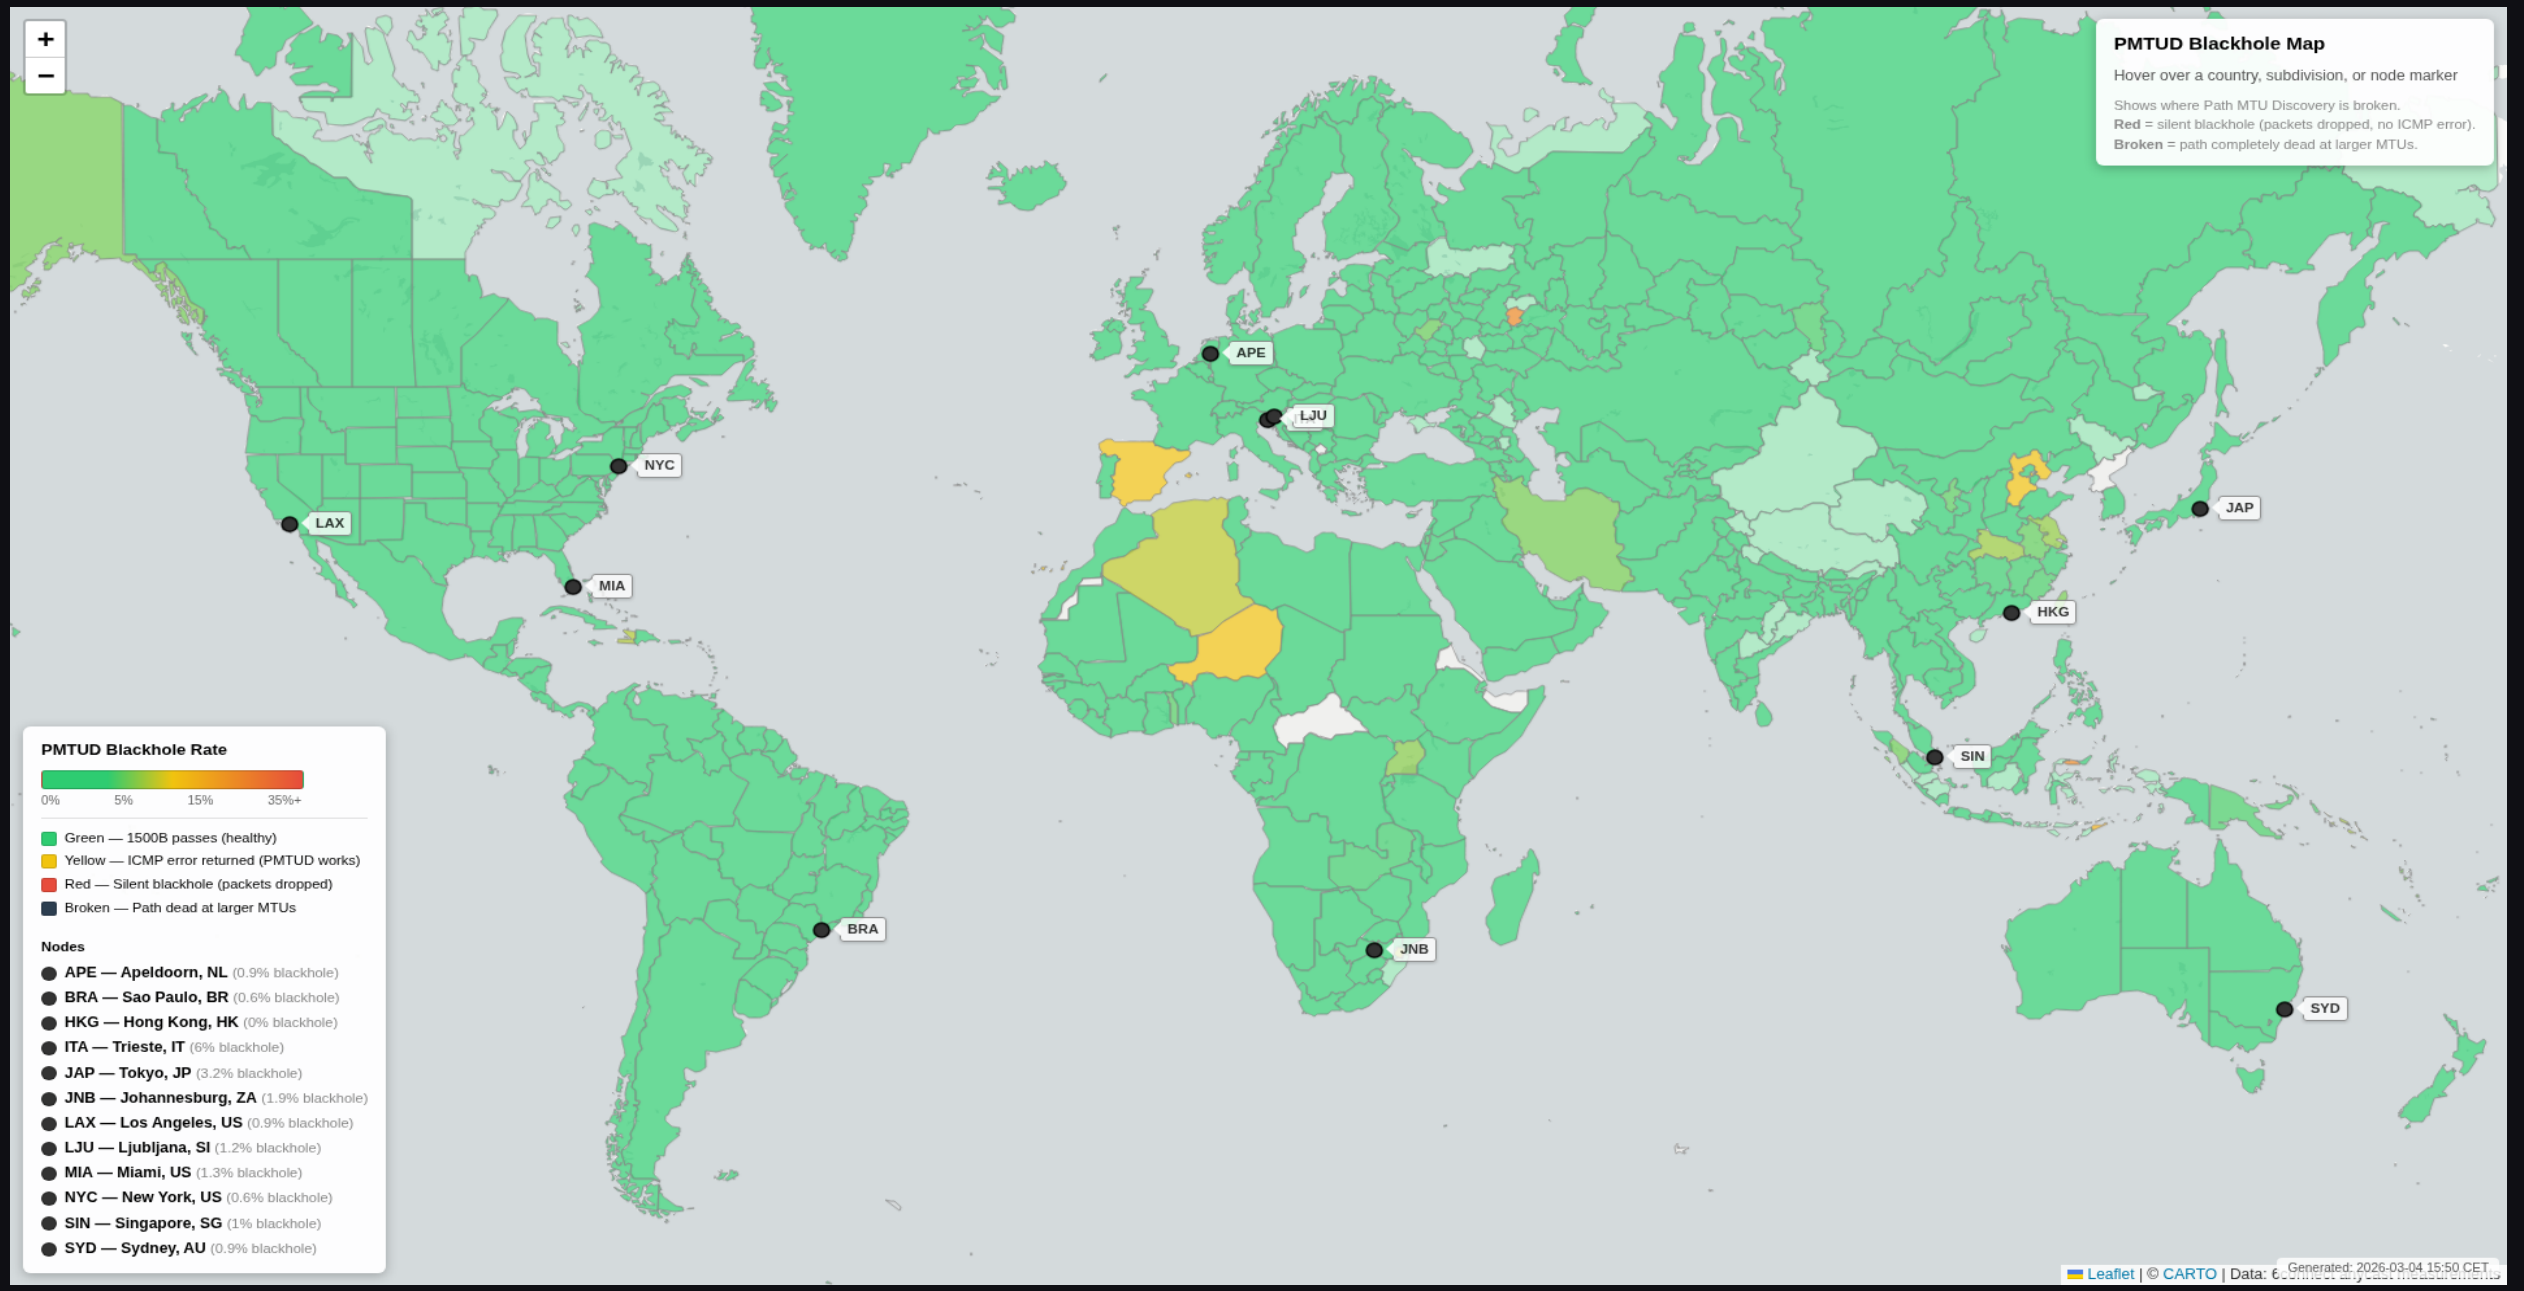

CHAPTER 4

Trust, but Verify

TCP SYN probing — is it really unreachable, or just ICMP-shy?

TCP SYN Probing

- Many hosts drop ICMP but respond to TCP
- yarrp TCP SYN to ports 80, 443, 53
- Anycast nodes capture SYN-ACK / RST via Go daemon
- 409,000 TCP responses in a single scan

The adventure of getting this to work:

- Step 1: Write raw TCP capture daemon. Easy.
- Step 2: Deploy to 33 nodes. Easy.
- Step 3: Get zero results. What?!
- Step 4: conntrack drops them as "INVALID"
- Step 5: Add NOTRACK rules. Success!

conntrack considers unsolicited SYN-ACK a threat. Fair point, honestly.

CHAPTER 5

Following the Breadcrumbs

Traceroutes — where exactly does it go wrong?

Targeted Traceroutes

- Only the top 50 anomalies — not everything
- dublin-traceroute for IPv4 (ECMP multipath)
- mtr for IPv6, 10 parallel traces from stats nodes

Classification pipeline:

- ICMP Shaping: TCP works, ICMP doesn't → not real
- Severe Loss: >15% packet loss → path issue
- User Congestion: high jitter, low loss → last-mile
- Routing anomaly → investigate!

dublin-traceroute writes to "trace.json" in current directory.
10 parallel = 10 race conditions. mktemp to the rescue.

CHAPTER 6

The View from Outside

BGP monitoring — what does the rest of the Internet see?

BGP Monitoring

Internal:

- 84 BIRD sessions across 33 nodes

External:

- RIPE RIS: 328 full-feed peers
- RouteViews: 3 collectors
- 100% visibility for all 6 prefixes

Real-time:

- RIS Live WebSocket stream
- Flap classification & email alerts

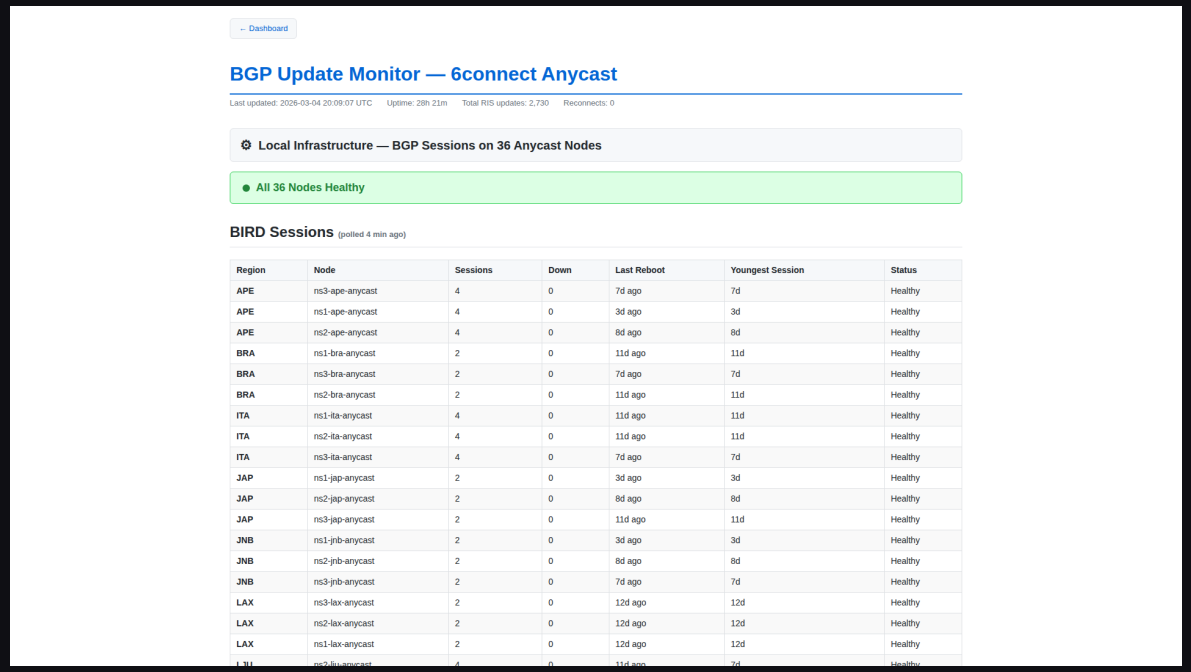

← Dashboard

BGP Update Monitor — 6connect Anycast

Last updated: 2026-03-04 20:09:07 UTC Uptime: 28h 21m Total RIS updates: 2,730 Reconnects: 0

⚙ Local Infrastructure — BGP Sessions on 36 Anycast Nodes

● All 36 Nodes Healthy

BIRD Sessions (polled 4 min ago)

Region	Node	Sessions	Down	Last Reboot	Youngest Session	Status
APE	ns3-ape-anycast	4	0	7d ago	7d	Healthy
APE	ns1-ape-anycast	4	0	3d ago	3d	Healthy
APE	ns2-ape-anycast	4	0	8d ago	8d	Healthy
BRA	ns1-bra-anycast	2	0	11d ago	11d	Healthy
BRA	ns3-bra-anycast	2	0	7d ago	7d	Healthy
BRA	ns2-bra-anycast	2	0	11d ago	11d	Healthy
ITA	ns1-ita-anycast	4	0	11d ago	11d	Healthy
ITA	ns2-ita-anycast	4	0	11d ago	11d	Healthy
ITA	ns3-ita-anycast	4	0	7d ago	7d	Healthy
JAP	ns1-jap-anycast	2	0	3d ago	3d	Healthy
JAP	ns2-jap-anycast	2	0	8d ago	8d	Healthy
JAP	ns3-jap-anycast	2	0	11d ago	11d	Healthy
JNB	ns1-jnb-anycast	2	0	3d ago	3d	Healthy
JNB	ns2-jnb-anycast	2	0	8d ago	8d	Healthy
JNB	ns3-jnb-anycast	2	0	7d ago	7d	Healthy
LAX	ns3-lax-anycast	2	0	12d ago	12d	Healthy
LAX	ns2-lax-anycast	2	0	12d ago	12d	Healthy
LAX	ns1-lax-anycast	2	0	12d ago	12d	Healthy
LJU	ns2-lju-anycast	4	0	11d ago	7d	Healthy

BGP Catchment — Community-Based Routing

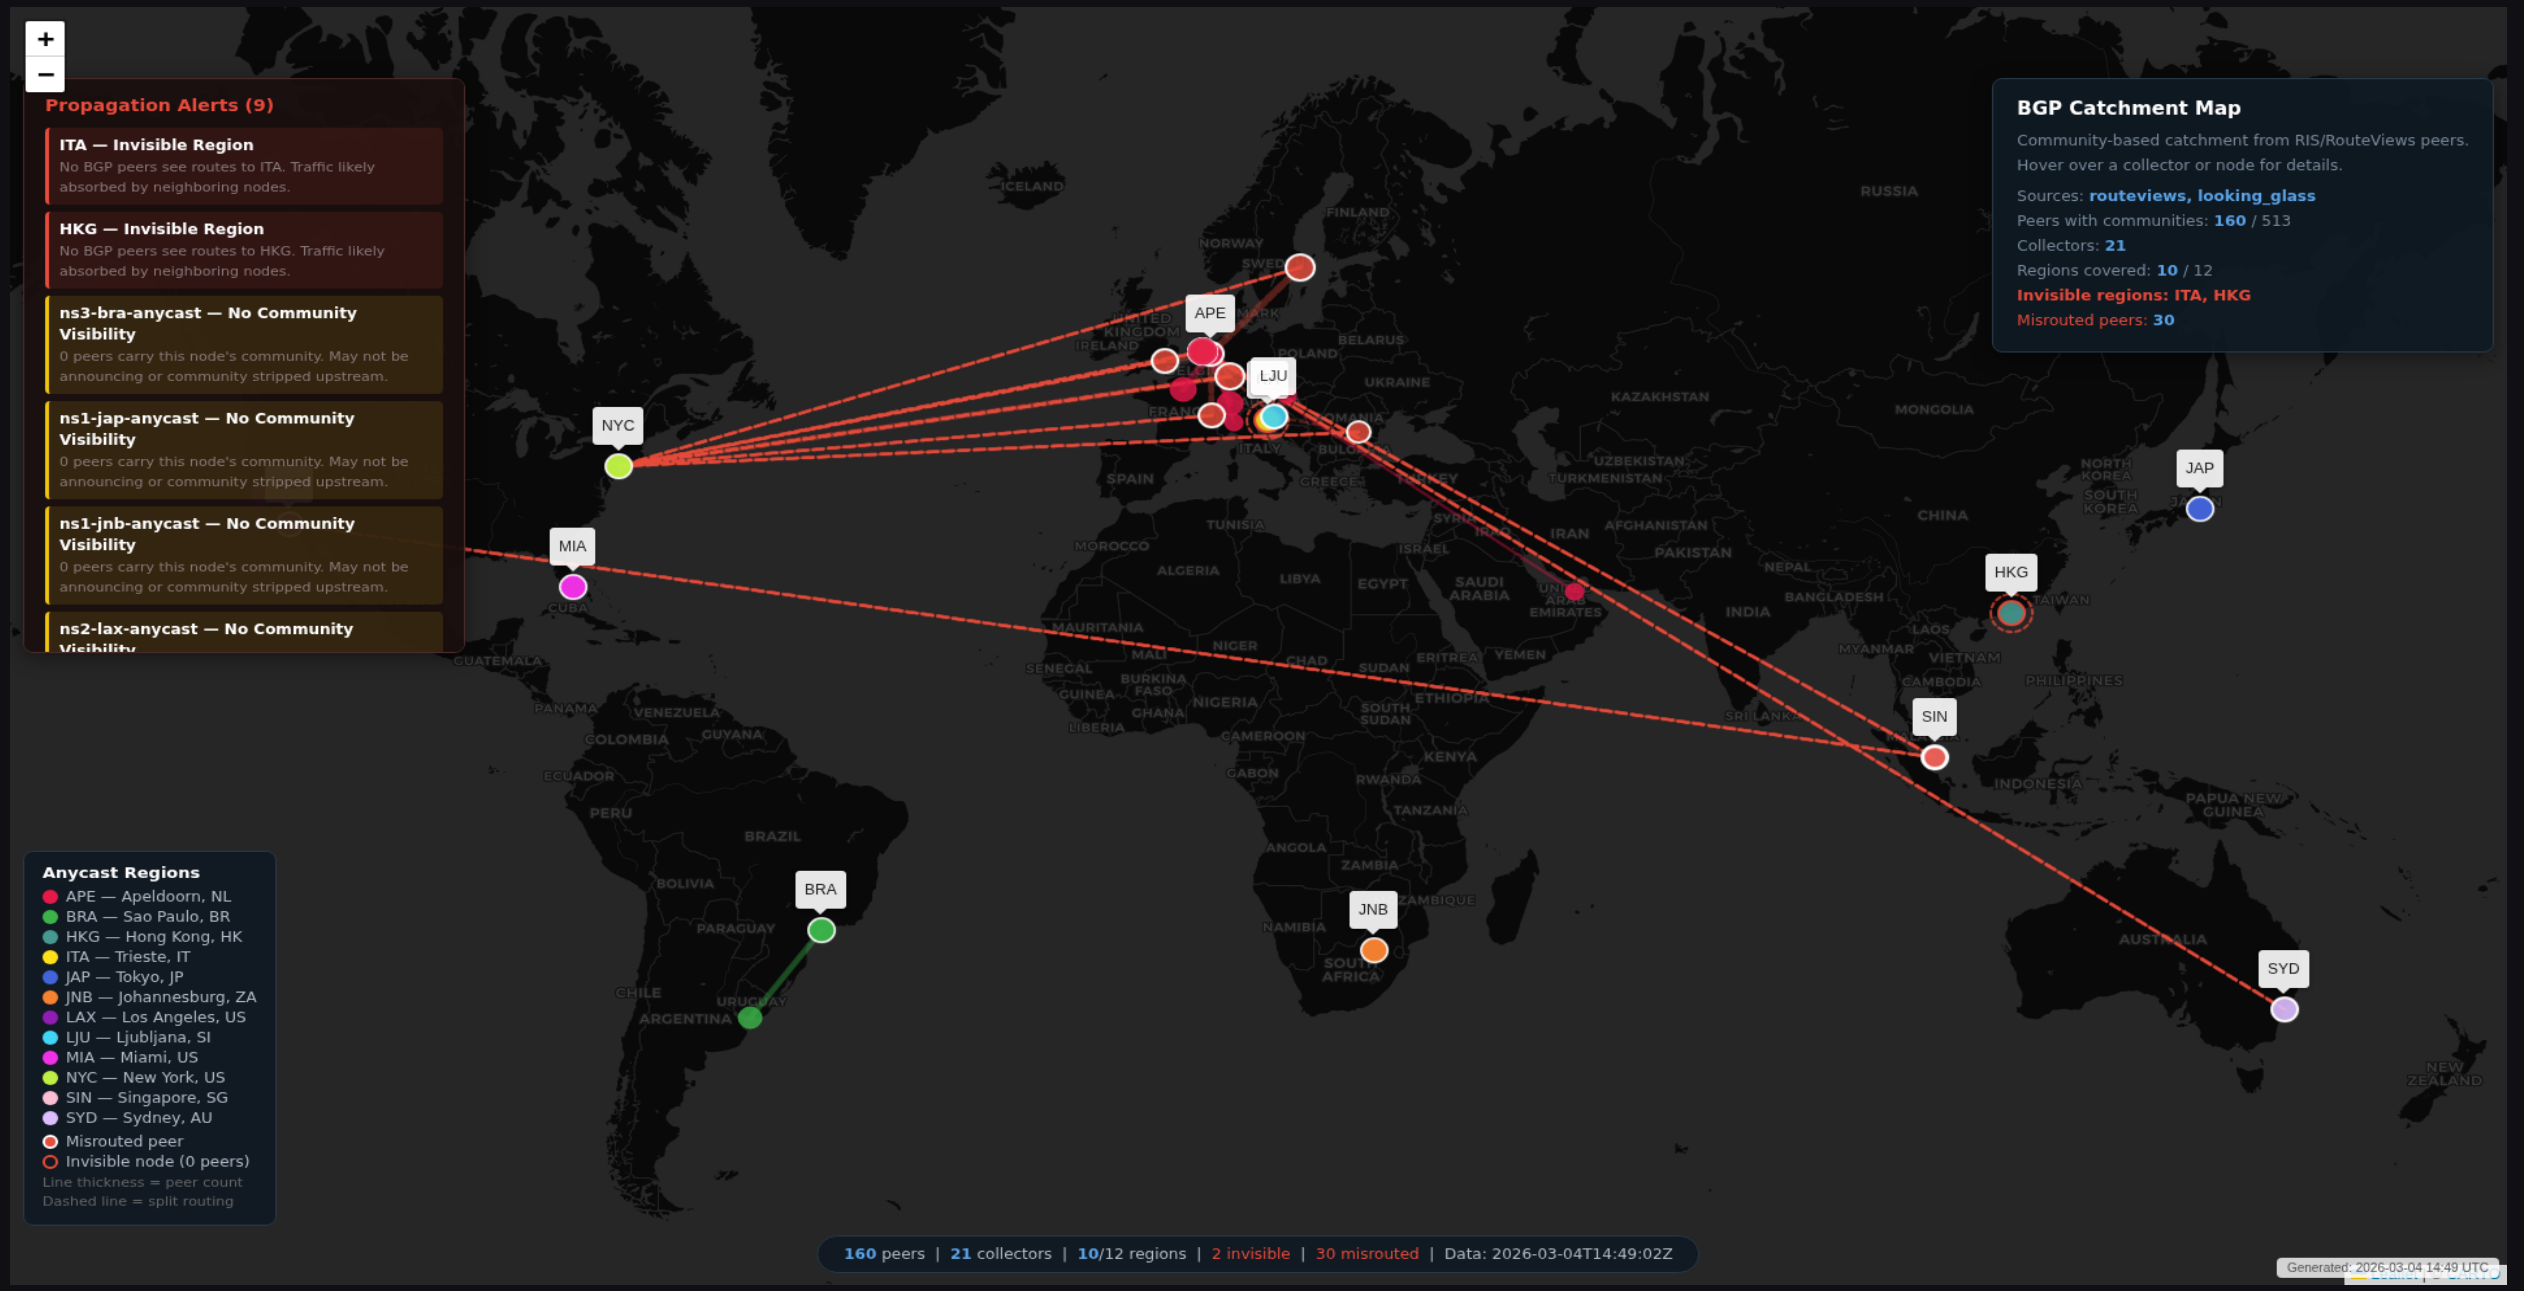

CHAPTER 7

The Bigger Picture

Correlating everything together

Connecting the Dots

- ICMP scan: this IP is routed to node X
- RTT: with 250ms latency and 15% loss
- PMTUD: 1500-byte packets silently dropped
- TCP: responds fine (SYN/ACK or RST) to SYN on port 443
- Traceroute: path goes through ISP Y in country Z
- BGP: ISP Y shouldn't route to node X at all

Each measurement alone is a data point.

Together, they tell a story.

Every day at 2am, the pipeline runs automatically.
By breakfast, we know exactly what happened overnight.

Geographic Misrouting — 10,000+ km Detours

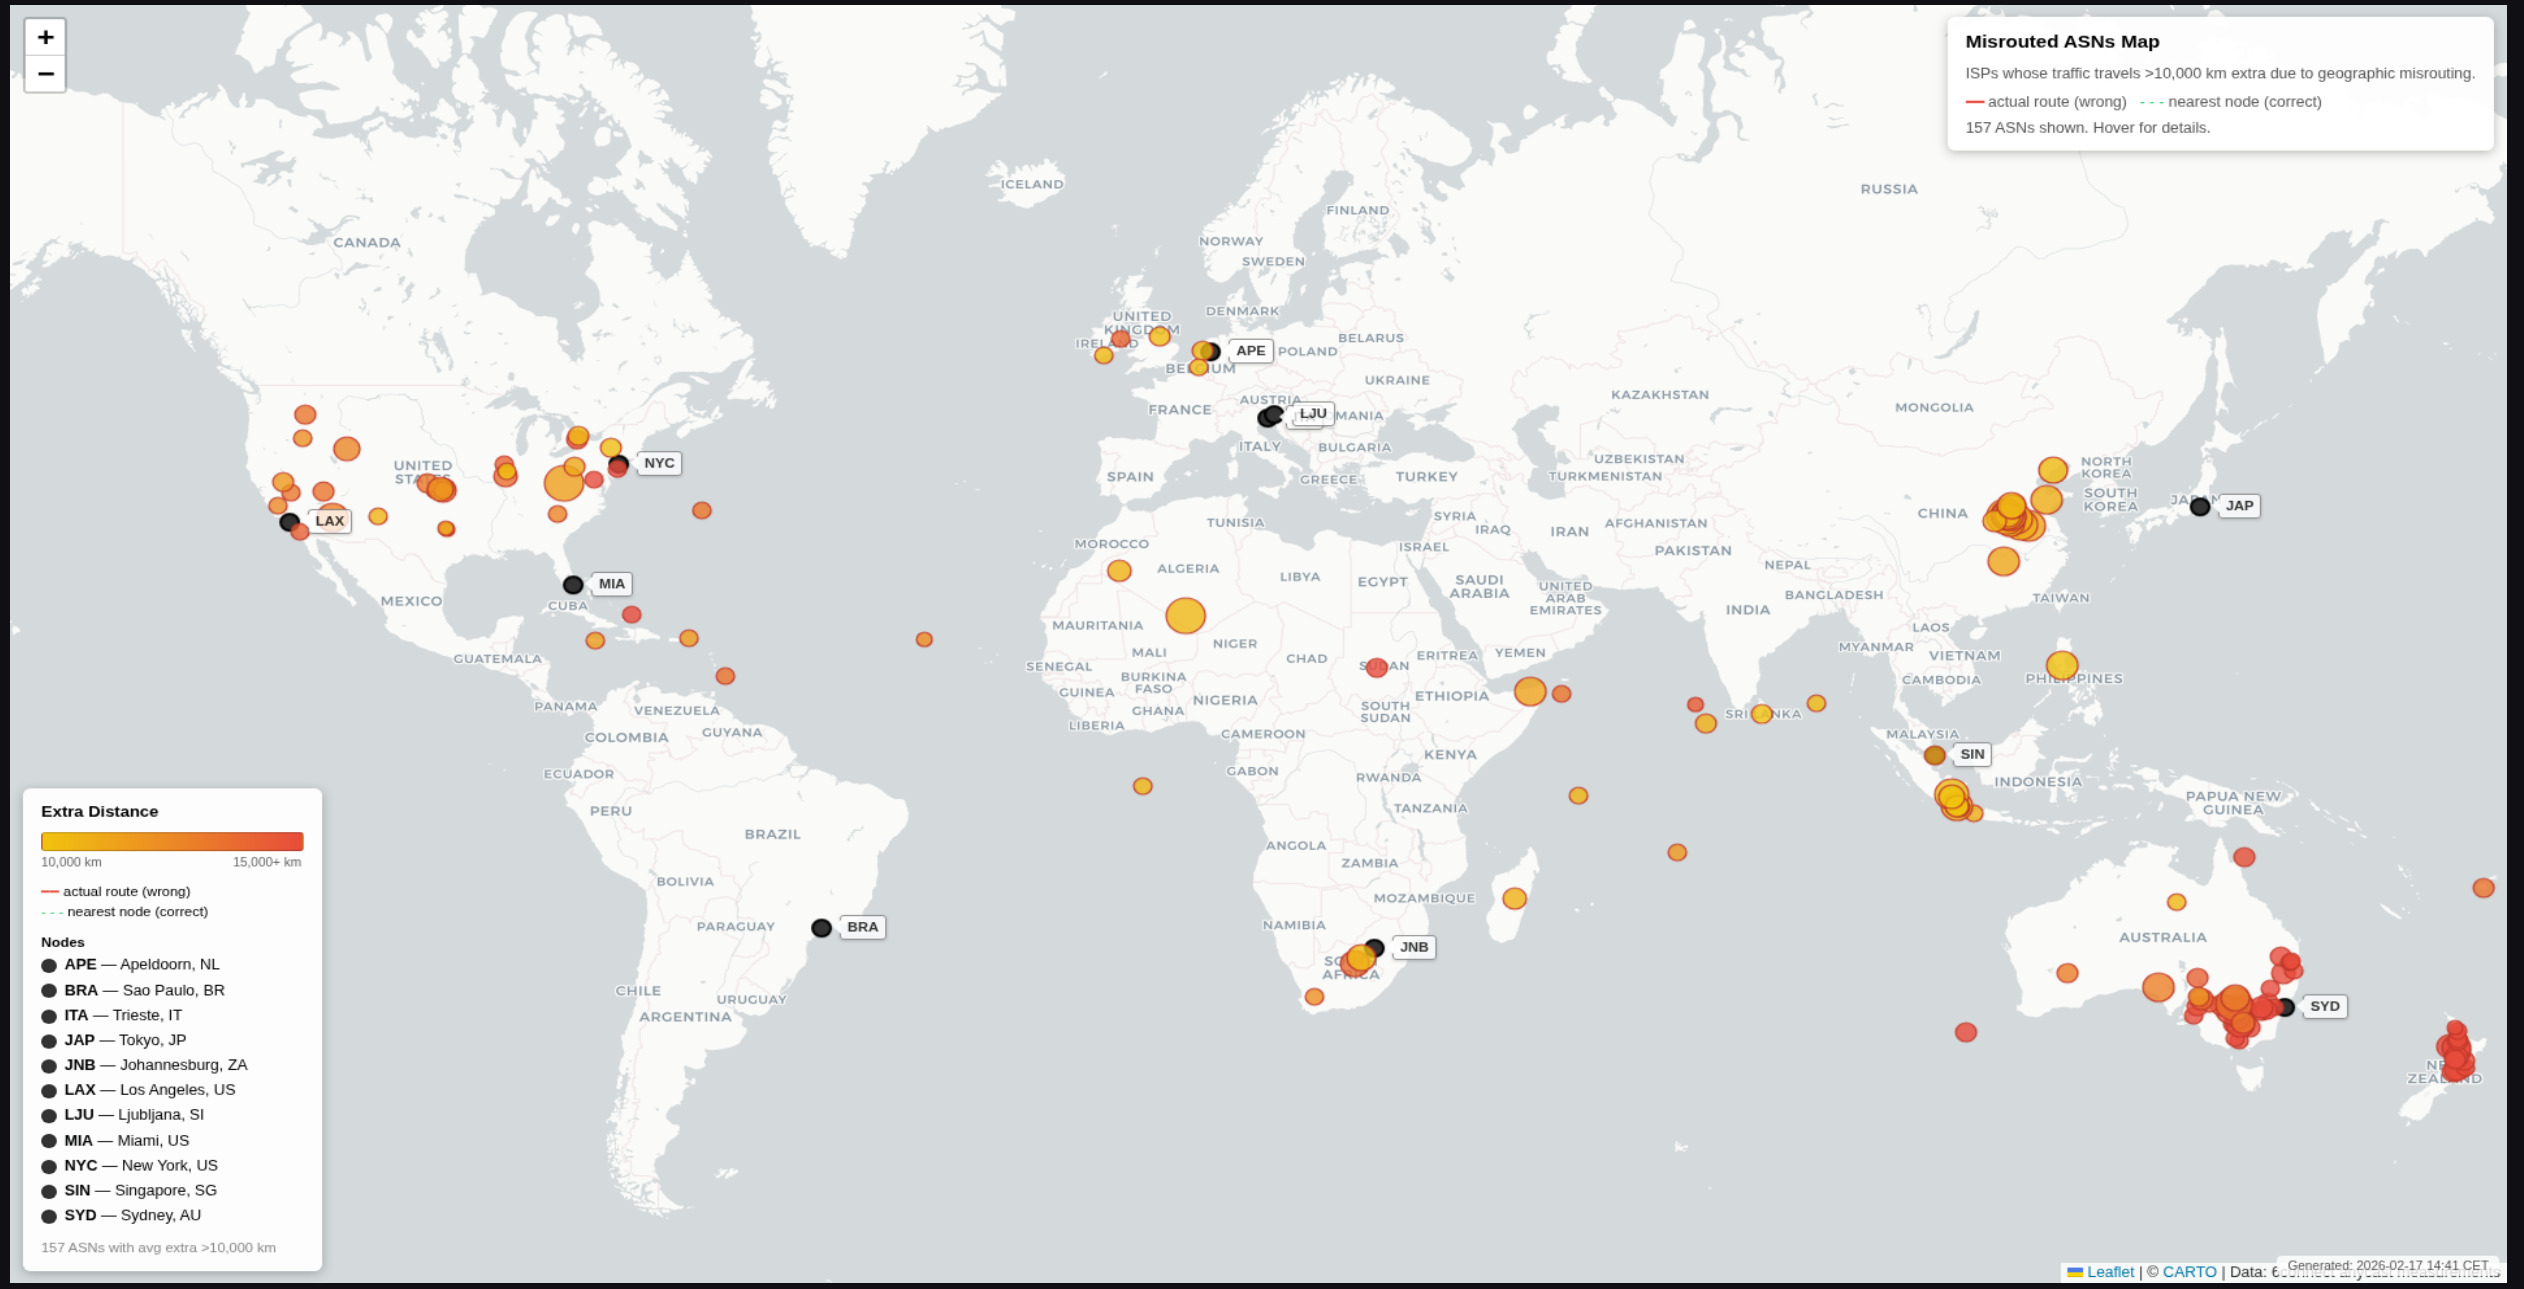

Problematic ASNs Report

Problematic ASNs

Anycast Discovery — ISPs with worst routing to the anycast network

← Back to Dashboard

193	32,853	65	107	36
PROBLEMATIC ASNS	AFFECTED IPS	RTT > 250MS	LOSS > 10%	BLACKHOLE > 25%

Worst RTT

338ms	331ms	324ms	323ms	315ms	305ms	304ms	297ms
AS131583	AS1659	AS136958	AS9930	AS33788	AS132199	AS17819	AS7539
WIZwireless Limite...	Taiwan Academic ...	China Unicom Gua...	TIME dotCom Berh...	KANARTEL (SD)	Globe Telecom Inc....	Equinix Asia Pacifi...	National Center fo...
11 IPs	22 IPs	82 IPs	315 IPs	11 IPs	218 IPs	26 IPs	31 IPs

296ms	295ms
AS37037	AS45433
Orange Madagasc...	Kappa Internet Ser...
43 IPs	14 IPs

Worst Packet Loss

43.3%	29.7%	29.0%	28.3%	27.5%	26.0%	26.0%	25.4%
AS9812	AS24138	AS137702	AS139462	AS210464	AS140553	AS4816	AS149178
Oriental Cable Net...	China TieTong Tele...	Nanjing, Jiangsu Pr...	CHINANET Guizho...	Andrew Kristuli (US)	CHINATELECOM XI...	China Telecom Gro...	China Telecom (CN)
18 IPs	33 IPs	58 IPs	12 IPs	16 IPs	10 IPs	60 IPs	22 IPs

24.7%	24.7%
AS134767	AS58541
CHINANET Sichuan...	Qingdao,266000 (...
17 IPs	47 IPs

Worst PMTUD Blackholes

100.0%	100.0%	100.0%	100.0%	94.4%	80.0%	80.0%	65.6%
AS210464	AS204490	AS6419	AS400519	AS44812	AS63961	AS21889	AS12732
Andrew Kristuli (US)	Kontel LLC (RU)	IDD Enterprises, L...	Developed Method...	Ip Server LLC (RU)	Bangladesh Resea...	Rapid Systems Cor...	GutCon GmbH (DE)
16 IPs	24 IPs	10 IPs	17 IPs	18 IPs	20 IPs	10 IPs	32 IPs

AI-Powered Analysis

- All data compiled into unified prompt (~200KB)
- Sent to Claude API every pipeline run
- Finds patterns in 775K rows that humans miss
- Recommends new node placements
- HTML report emailed automatically

13 analysis dimensions:

- Traceroutes, RTT, loss, PMTUD, BGP health
- Visibility, catchment, misrouting, stability

The AI also helped building this measurement system.
We're not sure if that's impressive or terrifying.

Marketing Numbers

99,000

probes/sec

1.7M

targets

775K

RTT measurements

409K

TCP responses

11

regions

33

nodes

~2h

full pipeline

Live Dashboard

Anycast Discovery

Distributed network measurement platform — 36 nodes across 12 regions

Catchment Map

Interactive world map showing which anycast node serves each country and US state. Hover for per-ISP routing breakdown.

INTERACTIVE

RTT Performance Map

Global latency heatmap colored by measured RTT. Green for fast, red for slow. Hover for per-region and per-ISP performance metrics.

INTERACTIVE

RTT History Map

Interactive time-slider map comparing RTT performance across measurement epochs. Slide between dates to see how latency evolves globally.

INTERACTIVE

PMTUD Blackhole Map

Path MTU Discovery blackhole distribution. Identifies countries and ISPs where large packets are silently dropped. Hover for per-ASN blackhole rates.

INTERACTIVE

Misrouted ASNs Map

Geographic misrouting visualization. Shows where ISP traffic reaches a far-away node instead of the nearest one. Hover for per-ASN misroute rates and extra distance.

INTERACTIVE

BGP Monitoring

Live BGP status dashboard. BIRD session health across all nodes, RIPE RIS and RouteViews visibility, flap detection with pivot analysis.

LIVE

Atlas Catchment Map

Anycast catchment from RIPE Atlas probes. DNS queries with NSID to ans1/ans2/ans3 over IPv4 and IPv6, showing which node serves each country.

Atlas RTT Map

DNS round-trip time from RIPE Atlas probes to anycast nodes. RTT gradient coloring with per-region and per-ISP performance breakdown.

What We Learned

- You can't measure anycast from inside the network
  - (well, you can, but you need to be clever about it)
- Every new measurement revealed unknown problems
- ICMP and TCP tell very different stories
- BGP "nearest" can be 10,000 km from geographic nearest
- Open-source tools have bugs. Kernels have bugs.
- The hardest part isn't collecting data —
  - it's knowing what to do with it

Thank You!

Questions?

jan@6connect.com sander@6connect.com

No anycast nodes were harmed in the making of this presentation.

A few kernel bugs were, though.