IPv6 operačně přestává být občan druhé kategorie

(doufejme)

Tomáš Hlaváček

Den IPv6 2026 · 4. 6. 2026

A public IPv6-only blockchain

Internet Computer (ICP / DFINITY)

- A public, decentralized blockchain by the DFINITY Foundation and a world-wide community — live since 2021.
- The network runs on IPv6 (-only) distributed consensus protocol that unites independent servers in commercial data centers around the world.
- Core blockchain protocol and most of the applications use exclusively IPv6 connectivity but each site has IPv4-enabled IC nodes → thousands of dual-stack targets, ideal for a large-scale v4-vs-v6 comparison.

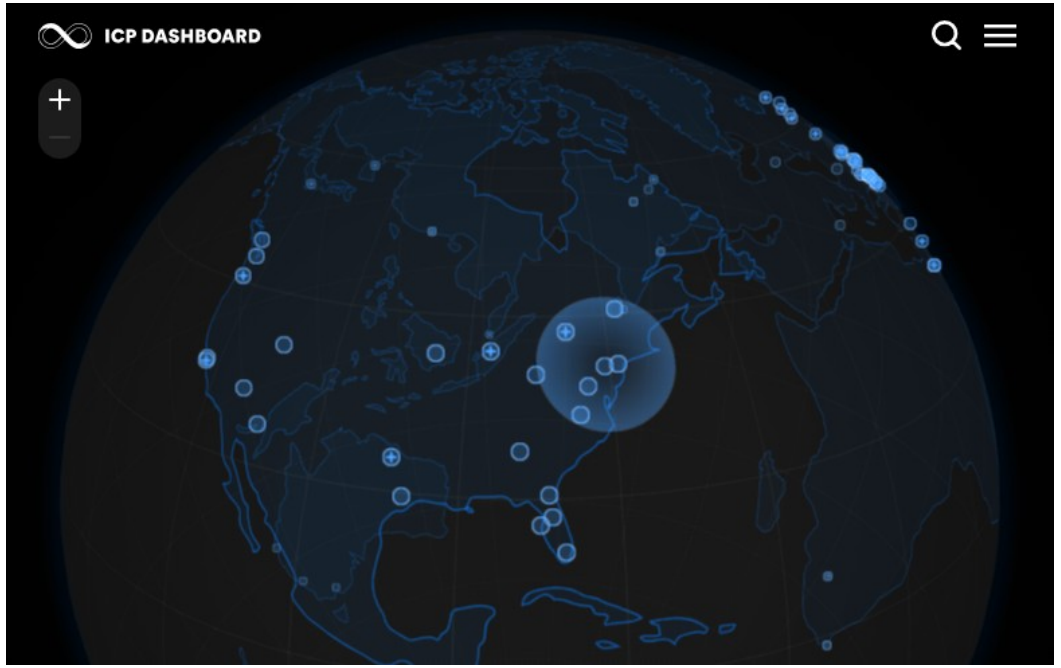

ICP data centers worldwide — the globe on the public ICP dashboard.
Source: dashboard.internetcomputer.org (public data), retrieved 2 Jun 2026

Data acquisition: targets from the public ICP dashboard

- **Target hosts and their IPv4/IPv6 addresses were scraped from the public ICP dashboard.**
- An independent infrastructure (bare-metal, AWS, RIPE Atlas) was used to conduct experiments – (no involvement of DFINITY).
- Measured with github.com/tmshlvck/teleping from 5 monitoring sites: 2x EU, 2x USA, 1x middle-east.
- ~6,000 monitored v4/v6 paths in total (target × site).

143	34	~1,450
data centers	countries	nodes

104	6
Autonomous Systems	continents

Active nodes in 34 countries — across all 6 continents:

Source: dashboard.internetcomputer.org (public data), retrieved 2 Jun 2026

What goes wrong: IPv6

- Path-MTU discovery breakage — Usually because ICMPv6 “Packet Too Big” gets filtered or is not emitted at all.
- Routing doesn’t clear up after an outage — stale v6 routes and neighbors linger long after v4 has recovered.
- L2/L3 multipath (LAG/ECMP) black-holes one shard — a bad bundle member silently drops part of the v6 flows.
- HW tables sync slower for v6 → transient black-holing.
- Broken end networks — misconfigured SLAAC and/or DHCPv6, wrong advertised PD validity times → renewals every 10s.
- systemd-networkd SLAAC-related bug(s) causing random v6 address and default route deconfig
- Apple devices won’t resolve AAAA over a v6-enabled tunnel/VPN interface when the physical interface has no v6.

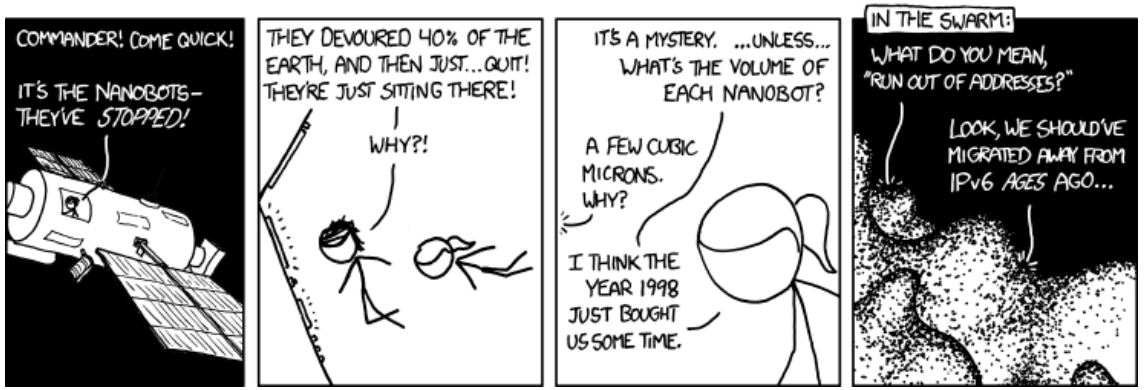

xkcd.com/865 “Nanobots” · CC BY-NC 2.5 · “...we should’ve migrated away from IPv6 ages ago...”

Dataset

Apr – Jun 2024

v6 slower to recover

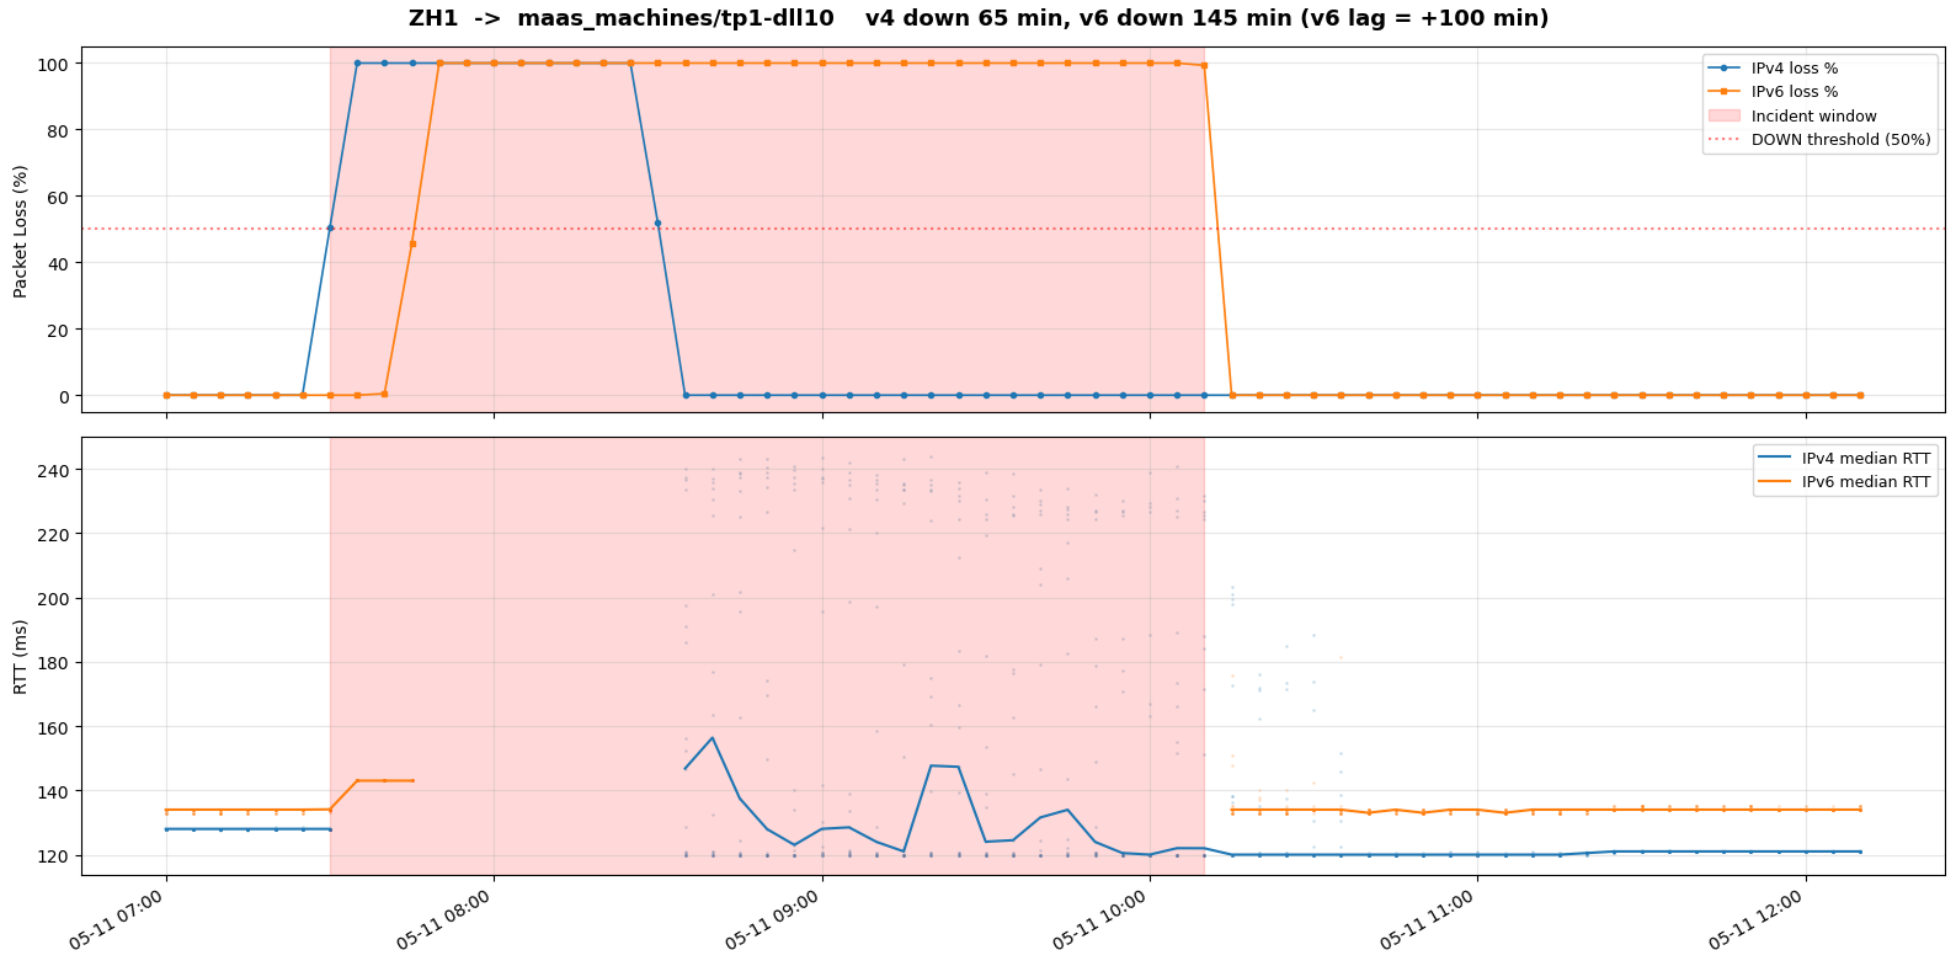

incident_00_lag_zh1_tp1-dll10

v6-only outage (v4 unaffected)

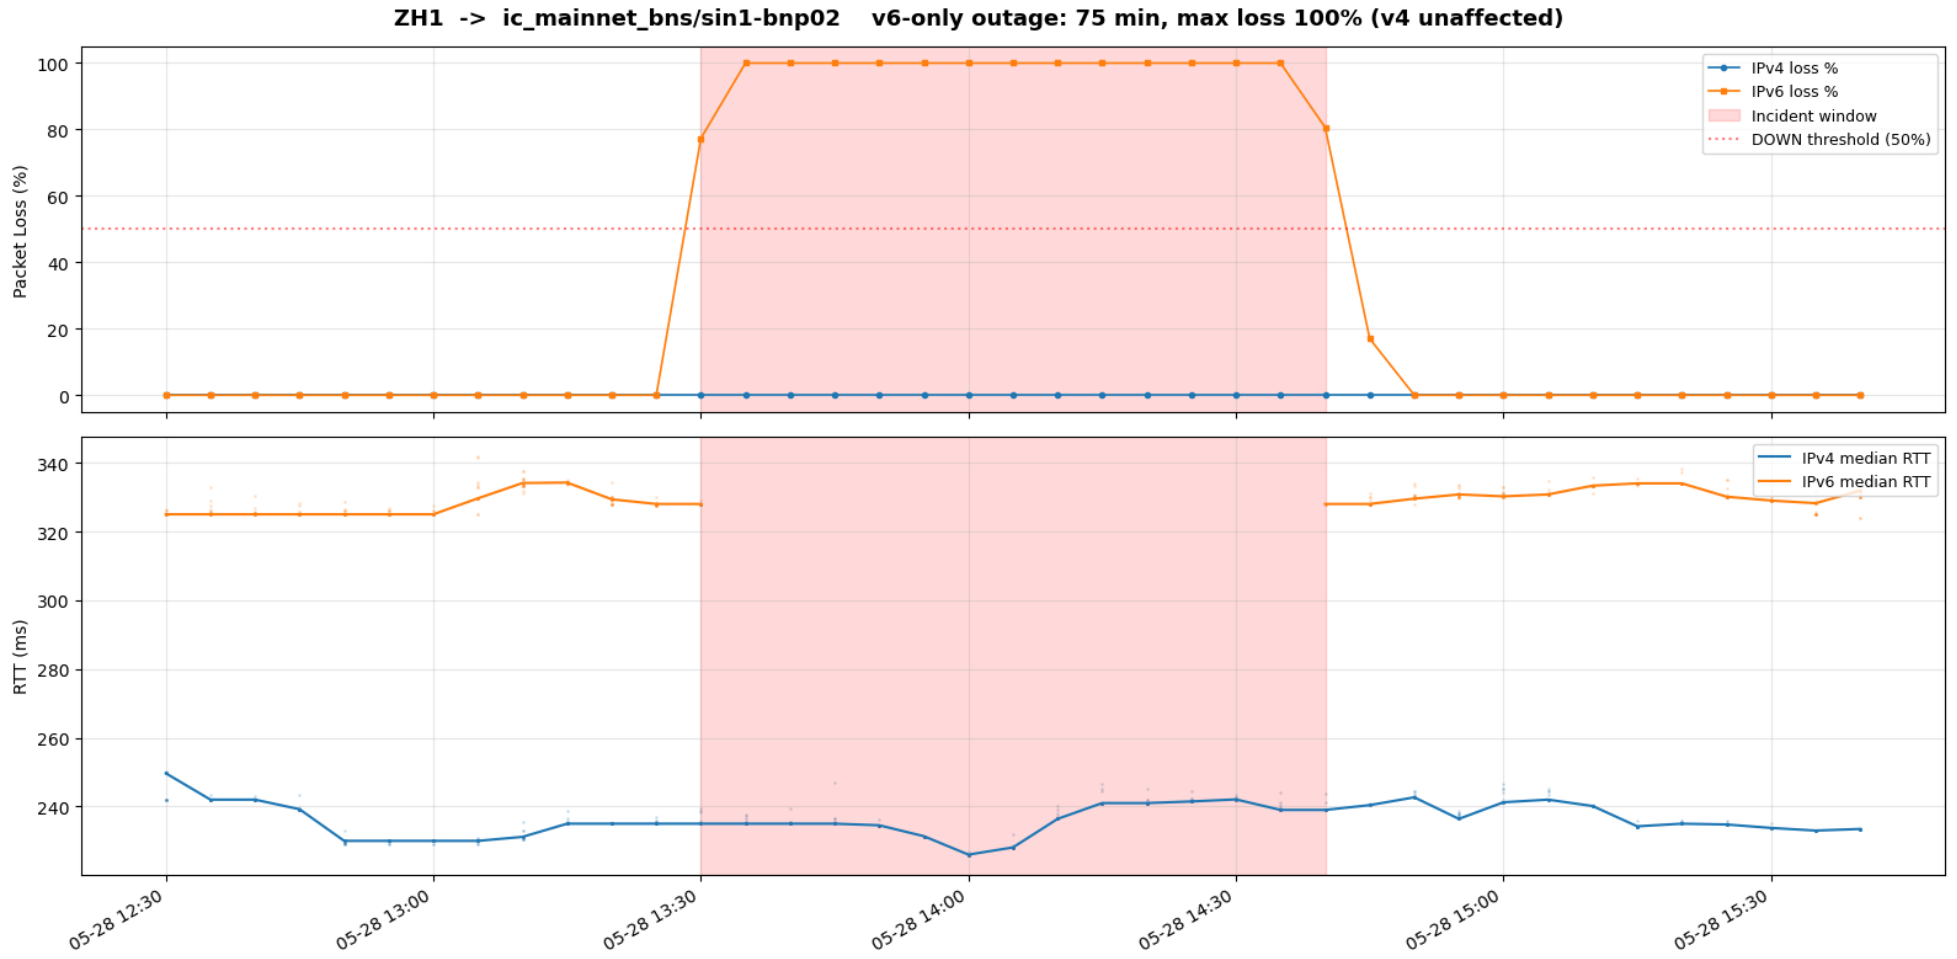

v6only_11_zh1_sin1-bnp02

Incident-rate comparison

```
Fraction of measurement intervals with ANY packet loss:
  IPv4 mean: 9.53%
  IPv6 mean: 14.81%
  Ratio (v6/v4): 1.55x

Fraction of intervals with TOTAL (100%) packet loss:
  IPv4 mean: 5.7699%
  IPv6 mean: 10.4041%
  Ratio (v6/v4): 1.80x

Per-target direction (any loss rate):
  v6 worse: 1237/1450 (85.3%)
  v4 worse: 38/1450 (2.6%)
  Equal:    175/1450
```

Distributions: loss / RTT / jitter / availability

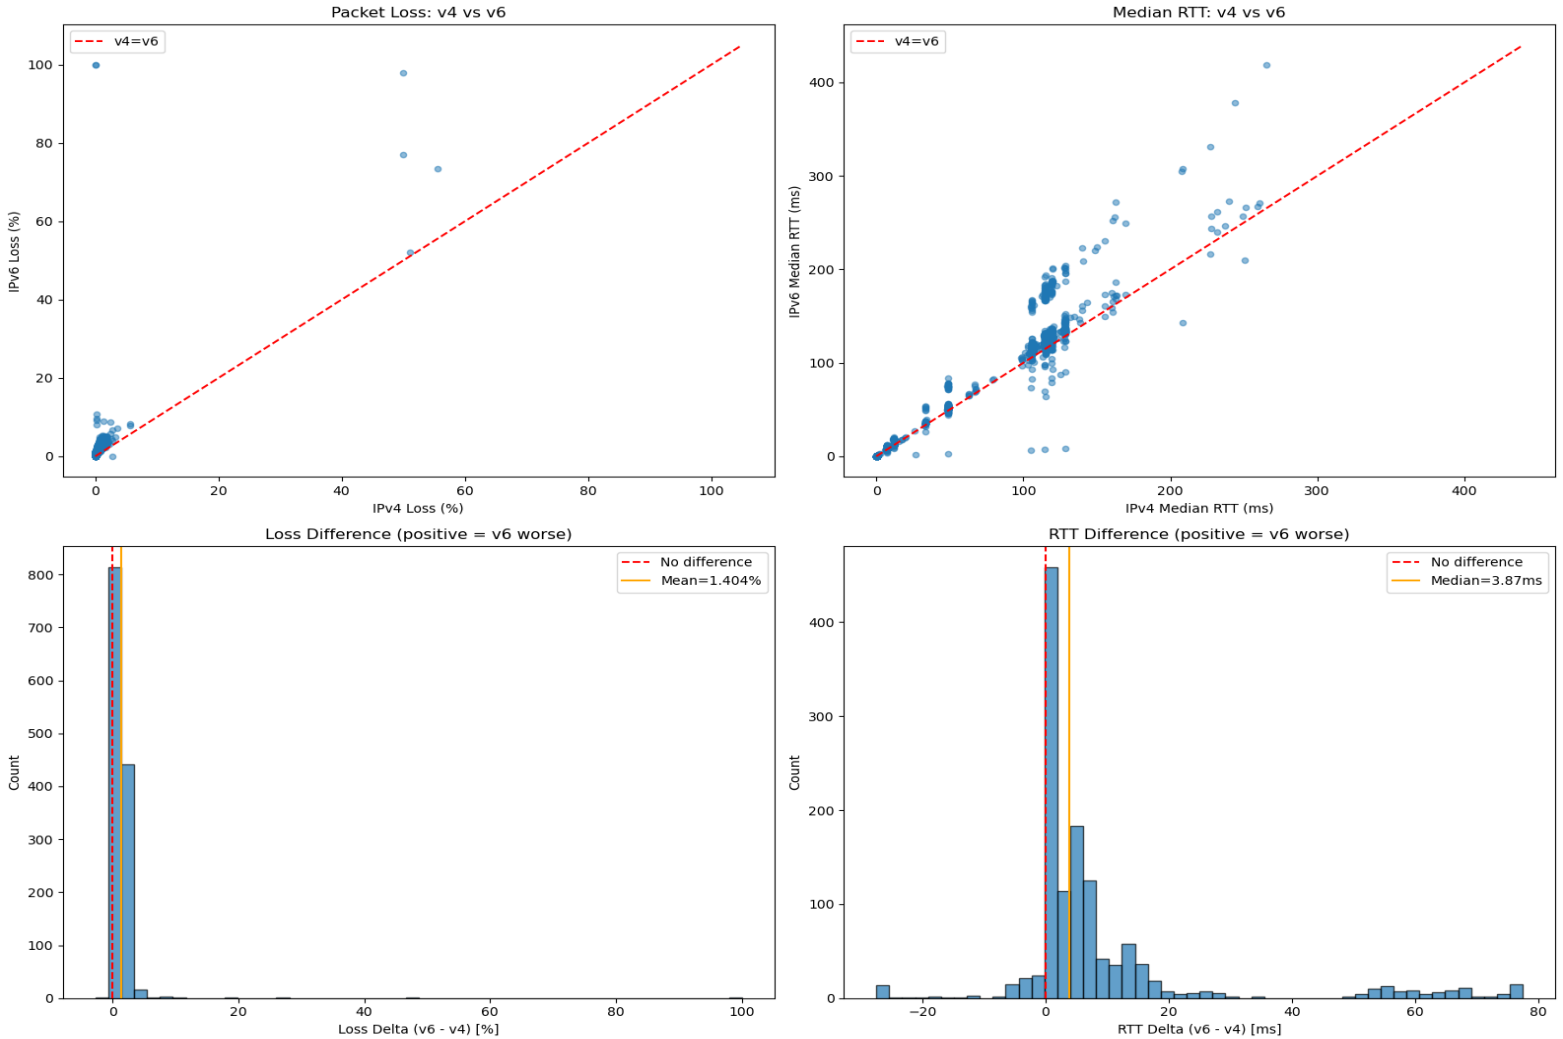

Overall IPv4 vs IPv6

	IPv4	IPv6	Delta (v6-v4)	Ratio (v6/v4)
Mean Loss (%)	0.5849	1.9892	1.4043	3.4009
Median Loss (%)	0.3165	1.4055	1.0890	4.4406
P95 Loss (%)	1.2997	4.0074	2.7077	3.0834
Mean Availability	0.9983	0.9962	-0.0021	0.9979
Mean Total-Loss Frac	0.0017	0.0038	0.0021	2.2574
Mean Median RTT (ms)	68.8509	77.2756	8.4246	1.1224
Median of Median RTT (ms)	103.3505	89.9647	-13.3858	0.8705
Mean Jitter (ms)	0.4308	0.4854	0.0546	1.1267
Median Jitter (ms)	0.2689	0.2941	0.0252	1.0938
Mean P95 RTT (ms)	75.1046	85.7243	10.6198	1.1414

Per-site box plots

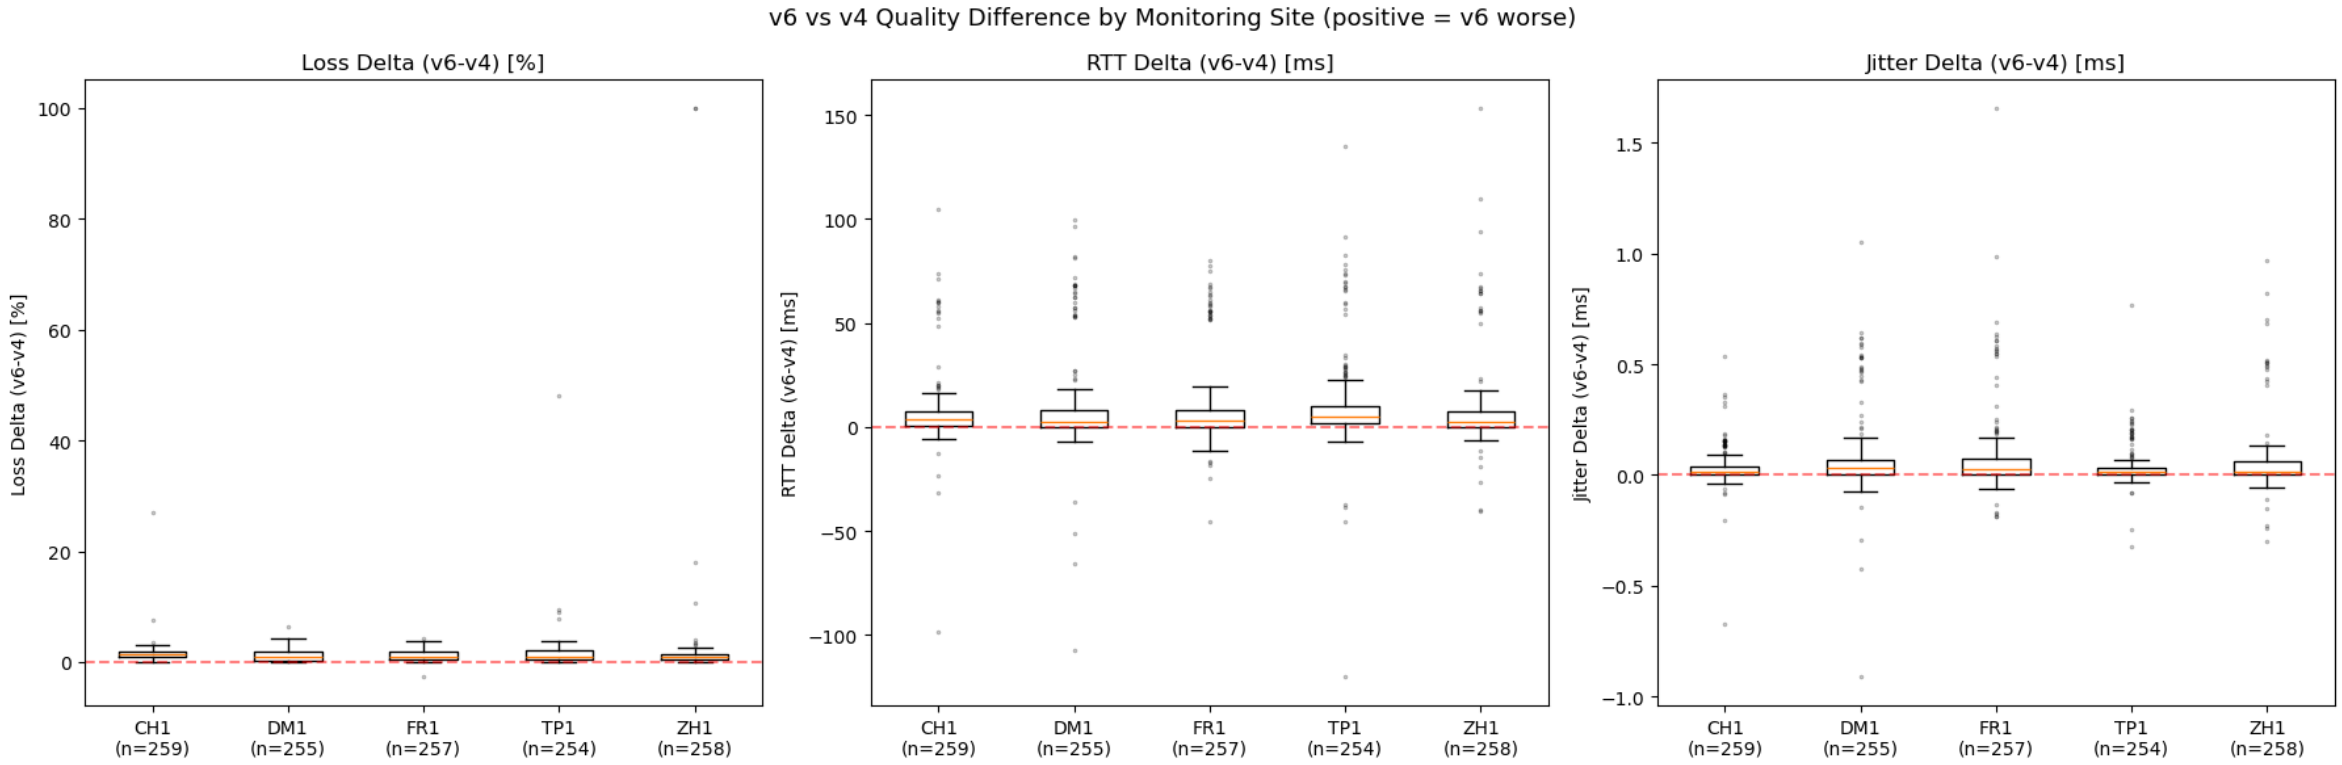

Dataset

Apr – Jun 2025

v6-only leading outage

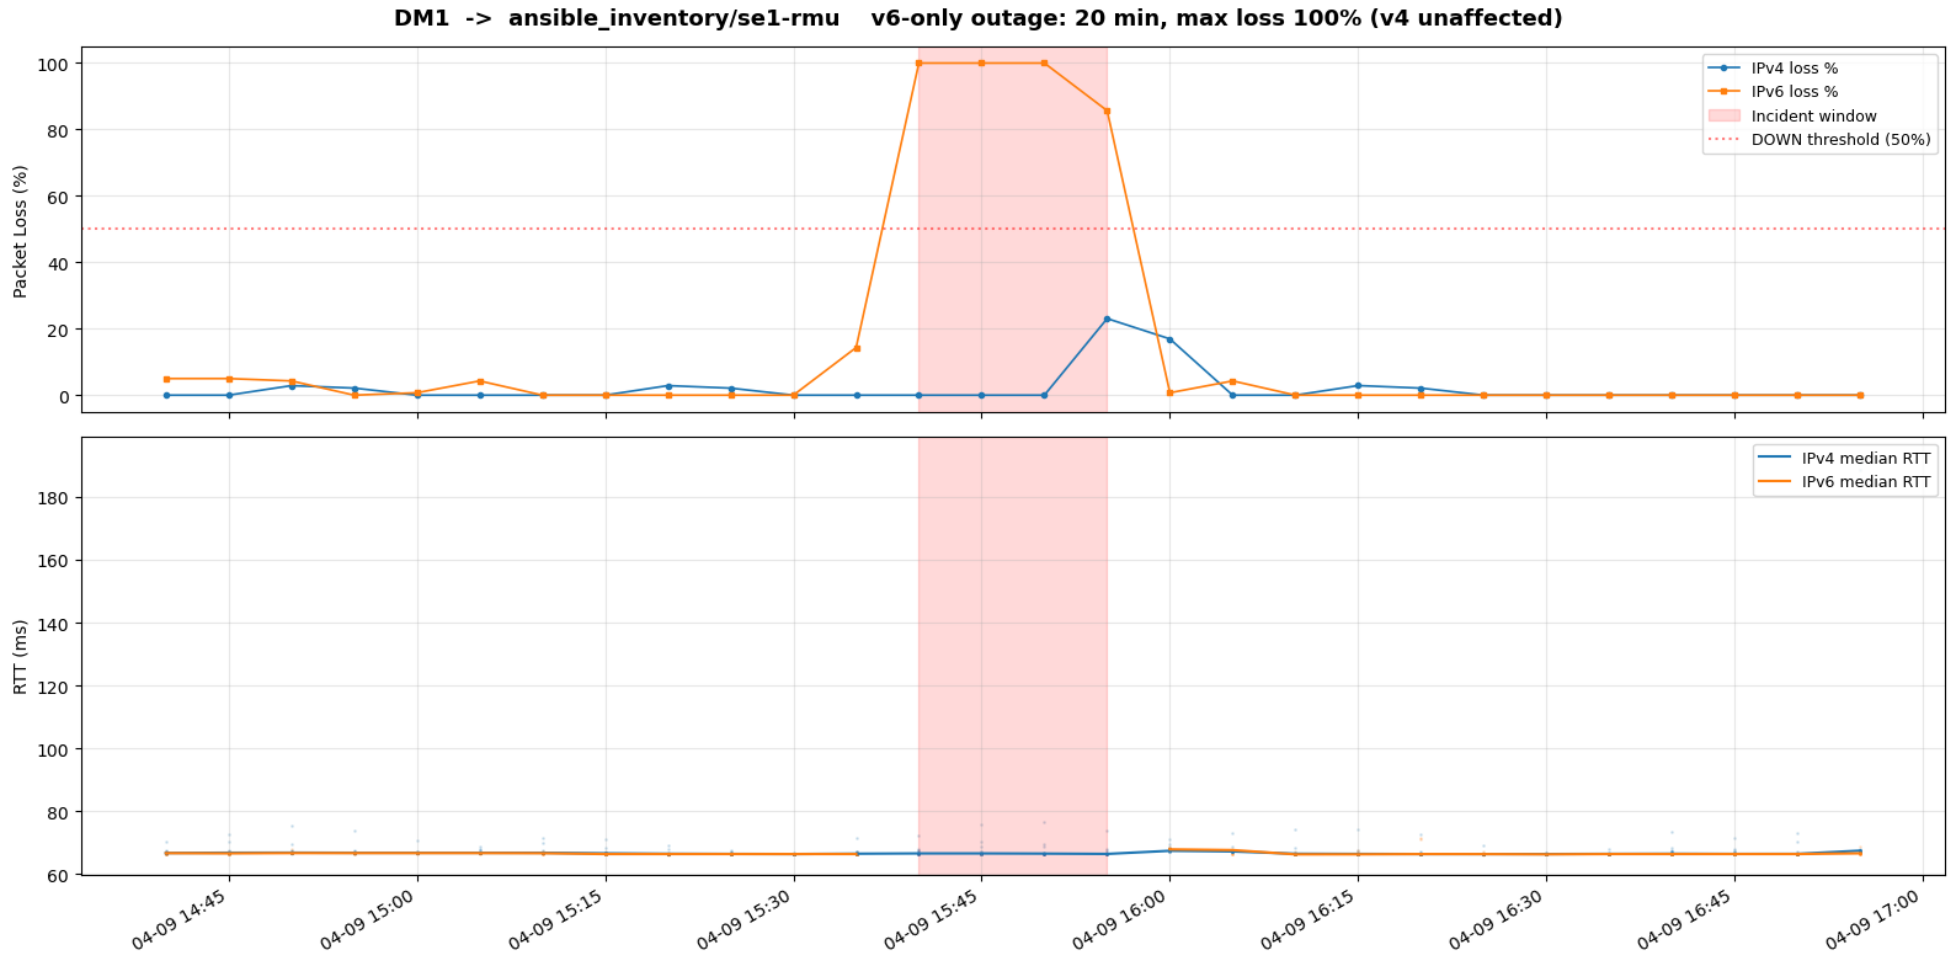

v6only_00_dm1_se1-rmu

Incident-rate comparison

```
Fraction of measurement intervals with ANY packet loss:
  IPv4 mean: 12.60%          (9.53%  in 2024)
  IPv6 mean: 19.75%          (14.81% in 2024)
  Ratio (v6/v4): 1.57x       (1.55x  in 2024)

Fraction of intervals with TOTAL (100%) packet loss:
  IPv4 mean: 6.9652%         (5.7699%  in 2024)
  IPv6 mean: 13.2594%        (10.4041% in 2024)
  Ratio (v6/v4): 1.90x       (1.80x    in 2024)

Per-target direction (any loss rate):
  v6 worse: 1387/1925 (72.1%)      (85.3% in 2024)
  v4 worse: 459/1925 (23.8%)       (2.6%  in 2024)
  Equal:    79/1925
```

Distributions: loss / RTT / jitter / availability

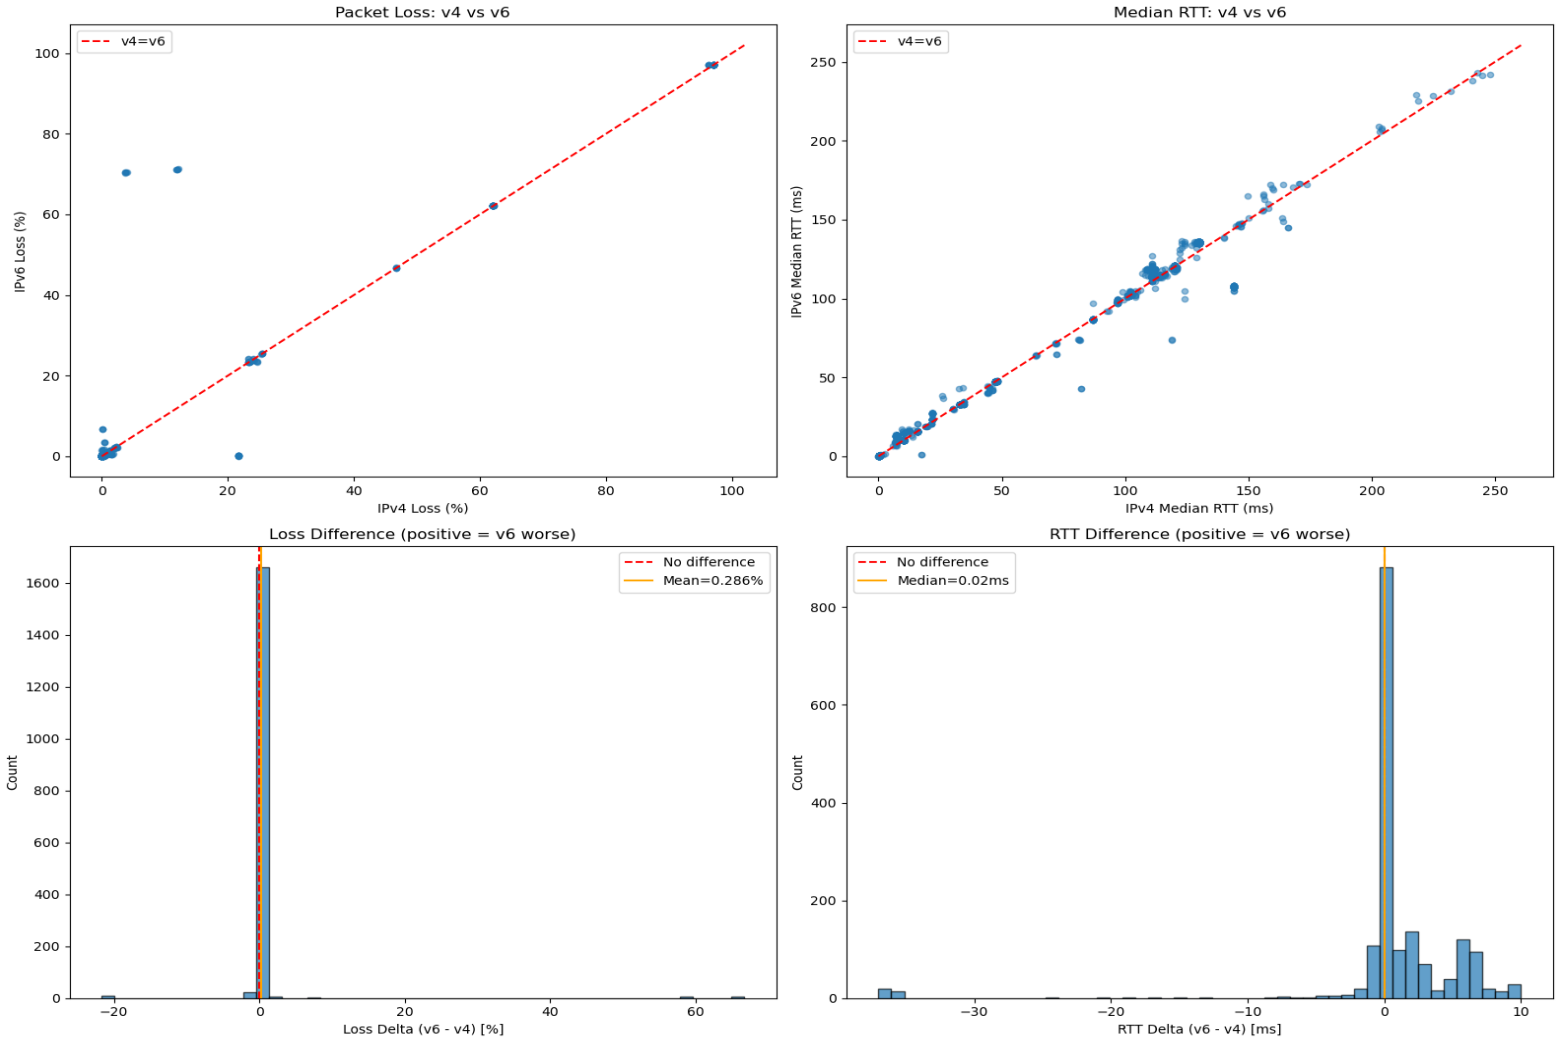

Overall IPv4 vs IPv6

	IPv4	IPv6	Delta (v6-v4)	Ratio (v6/v4)
Mean Loss (%)	3.8237	4.1102	0.2865	1.0749
Median Loss (%)	0.1820	0.2031	0.0211	1.1161
P95 Loss (%)	23.2147	23.5886	0.3739	1.0161
Mean Availability	0.9629	0.9602	-0.0027	0.9972
Mean Total-Loss Frac	0.0371	0.0398	0.0027	1.0717
Mean Median RTT (ms)	66.0906	66.7351	0.6446	1.0098
Median of Median RTT (ms)	80.8917	71.7750	-9.1167	0.8873
Mean Jitter (ms)	0.1257	0.0608	-0.0649	0.4836
Median Jitter (ms)	0.1018	0.0446	-0.0572	0.4383
Mean P95 RTT (ms)	69.9549	75.9411	5.9862	1.0856

Per-site box plots

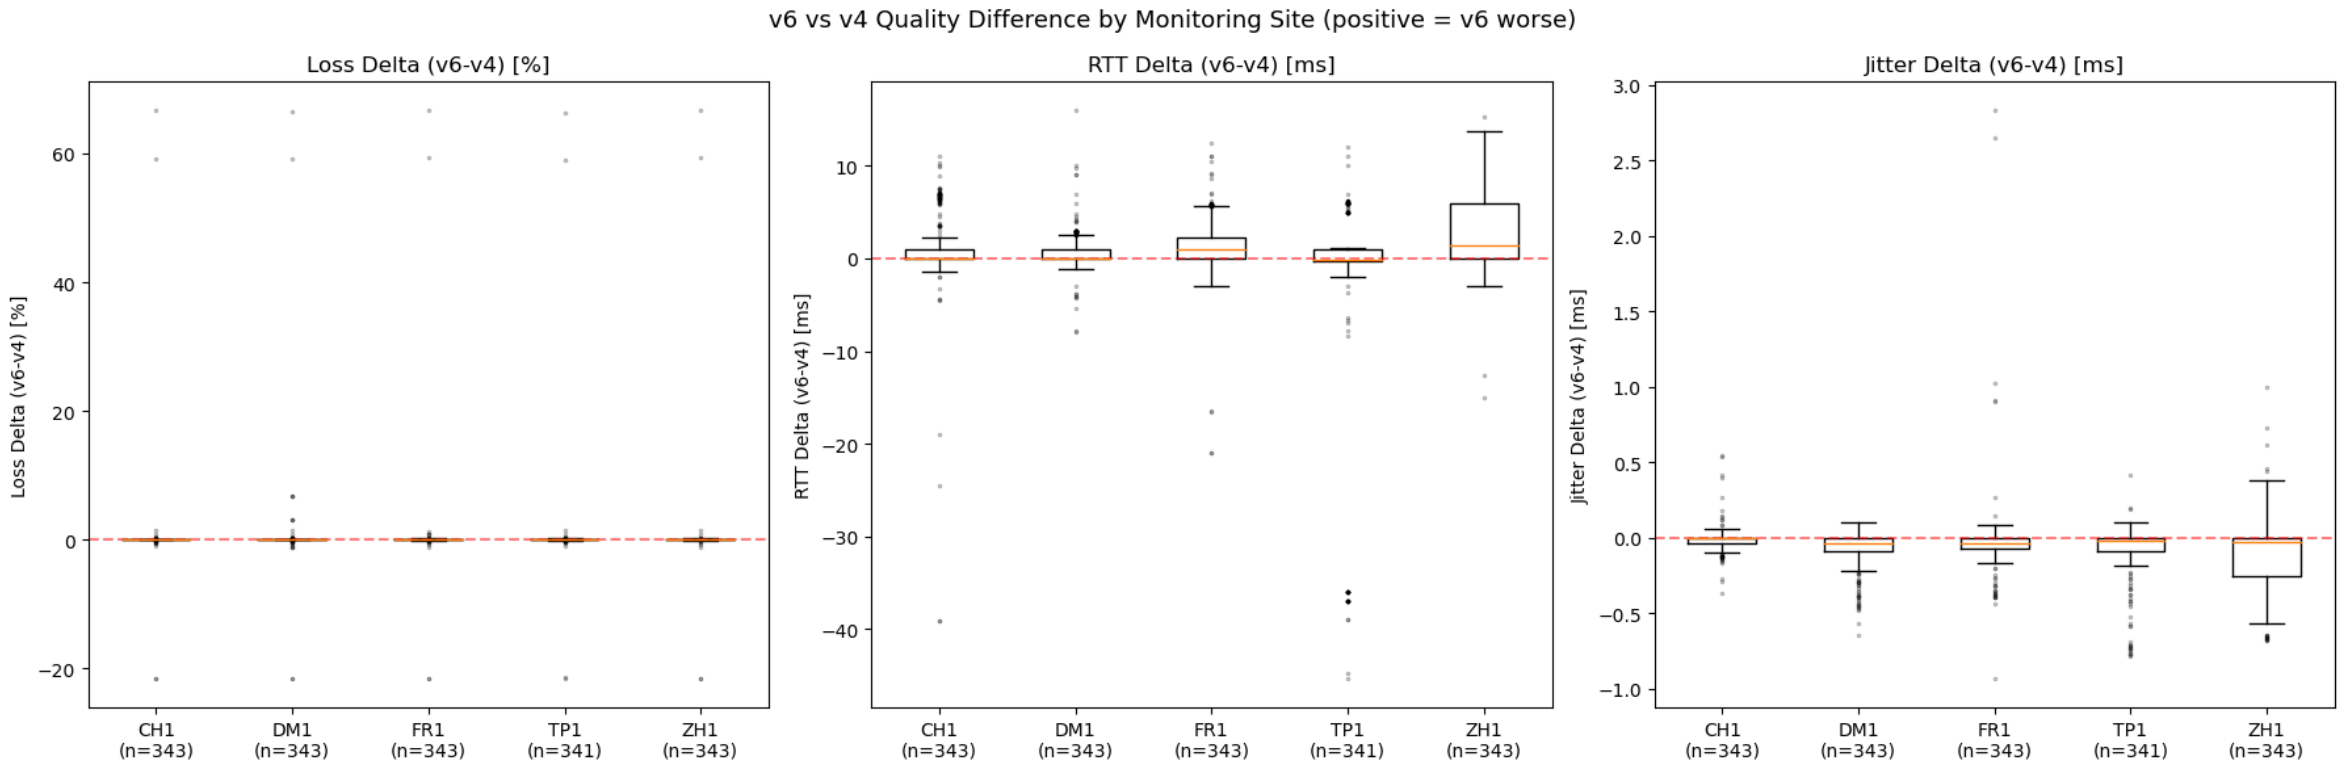

Dataset

Mar – May 2026

Incident-rate comparison

```
Fraction of measurement intervals with ANY packet loss:
  IPv4 mean: 9.70%          (12.60% in 2025)
  IPv6 mean: 14.30%         (19.75% in 2025)
  Ratio (v6/v4): 1.47x      (1.57x  in 2025)

Fraction of intervals with TOTAL (100%) packet loss:
  IPv4 mean: 8.9438%        (6.9652%  in 2025)
  IPv6 mean: 13.2058%       (13.2594% in 2025)
  Ratio (v6/v4): 1.48x      (1.90x    in 2025)

Per-target direction (any loss rate):
  v6 worse: 791/1455 (54.4%)   (72.1% in 2025)
  v4 worse: 374/1455 (25.7%)   (23.8% in 2025)
  Equal:    290/1455
```

Distributions: loss / RTT / jitter / availability

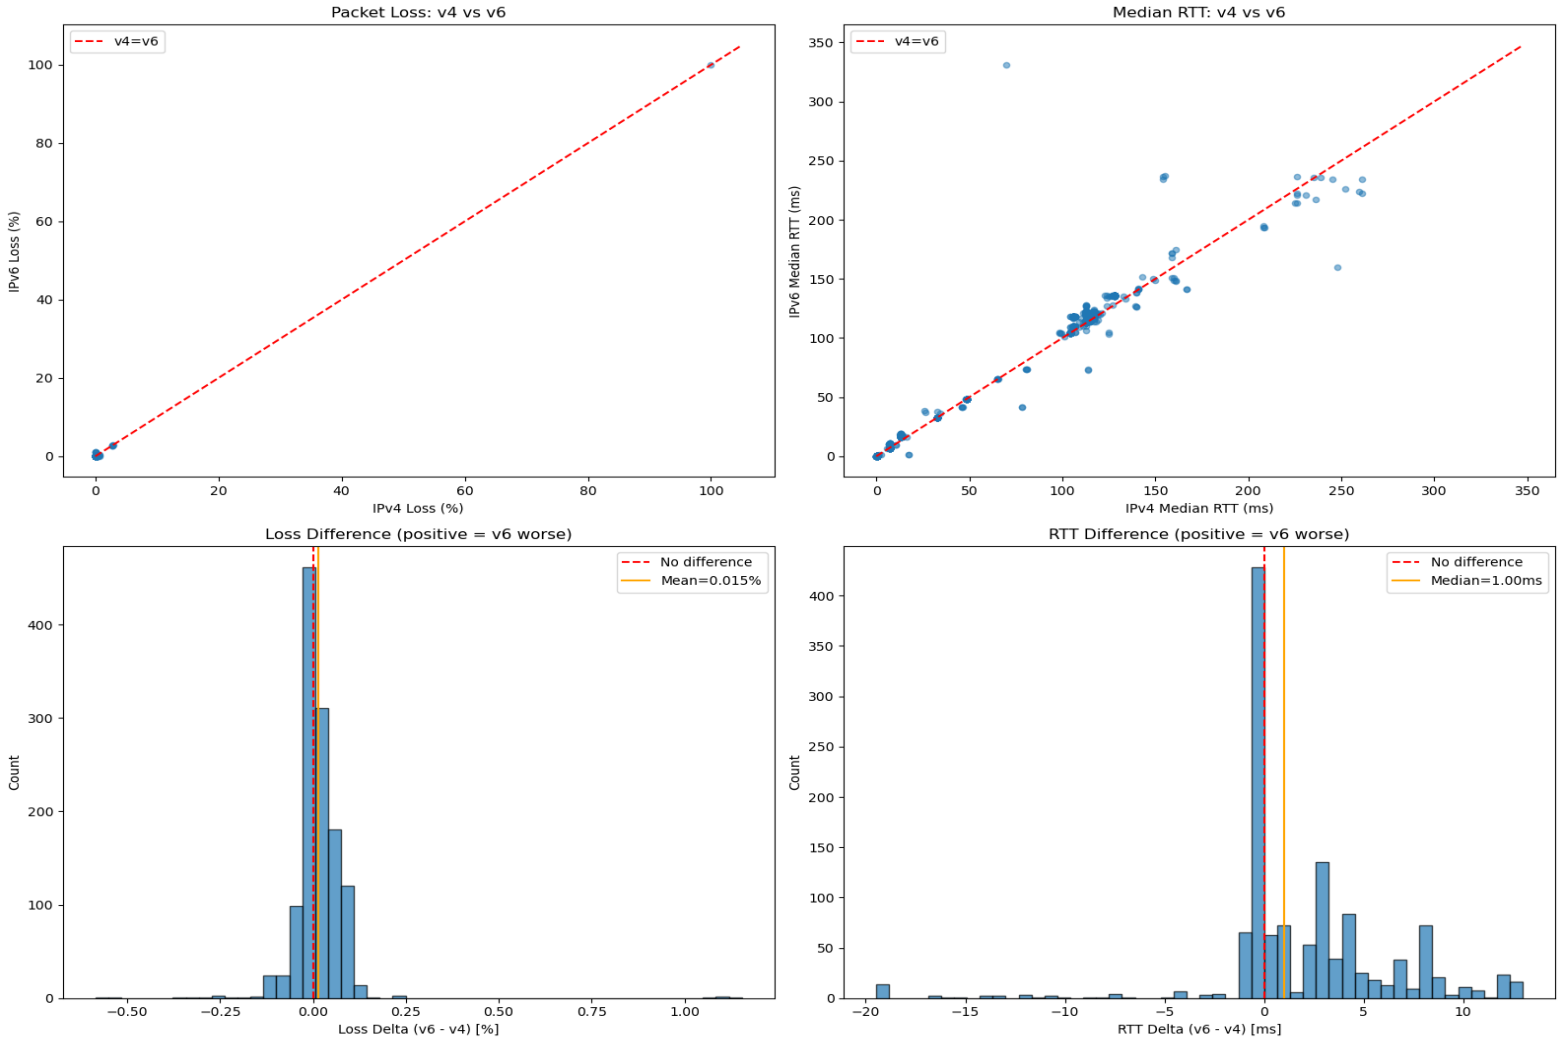

Overall IPv4 vs IPv6

	IPv4	IPv6	Delta (v6-v4)	Ratio (v6/v4)
Mean Loss (%)	0.1425	0.1577	0.0152	1.1067
Median Loss (%)	0.0290	0.0483	0.0193	1.6673
P95 Loss (%)	0.1698	0.1786	0.0089	1.0521
Mean Availability	0.9991	0.9991	-0.0000	1.0000
Mean Total-Loss Frac	0.0009	0.0009	0.0000	1.0015
Mean Median RTT (ms)	68.4490	70.6772	2.2283	1.0326
Median of Median RTT (ms)	104.0000	104.0000	0.0000	1.0000
Mean Jitter (ms)	0.1560	0.0484	-0.1076	0.3102
Median Jitter (ms)	0.1032	0.0102	-0.0929	0.0994
Mean P95 RTT (ms)	71.2558	74.3189	3.0631	1.0430

Per-site box plots

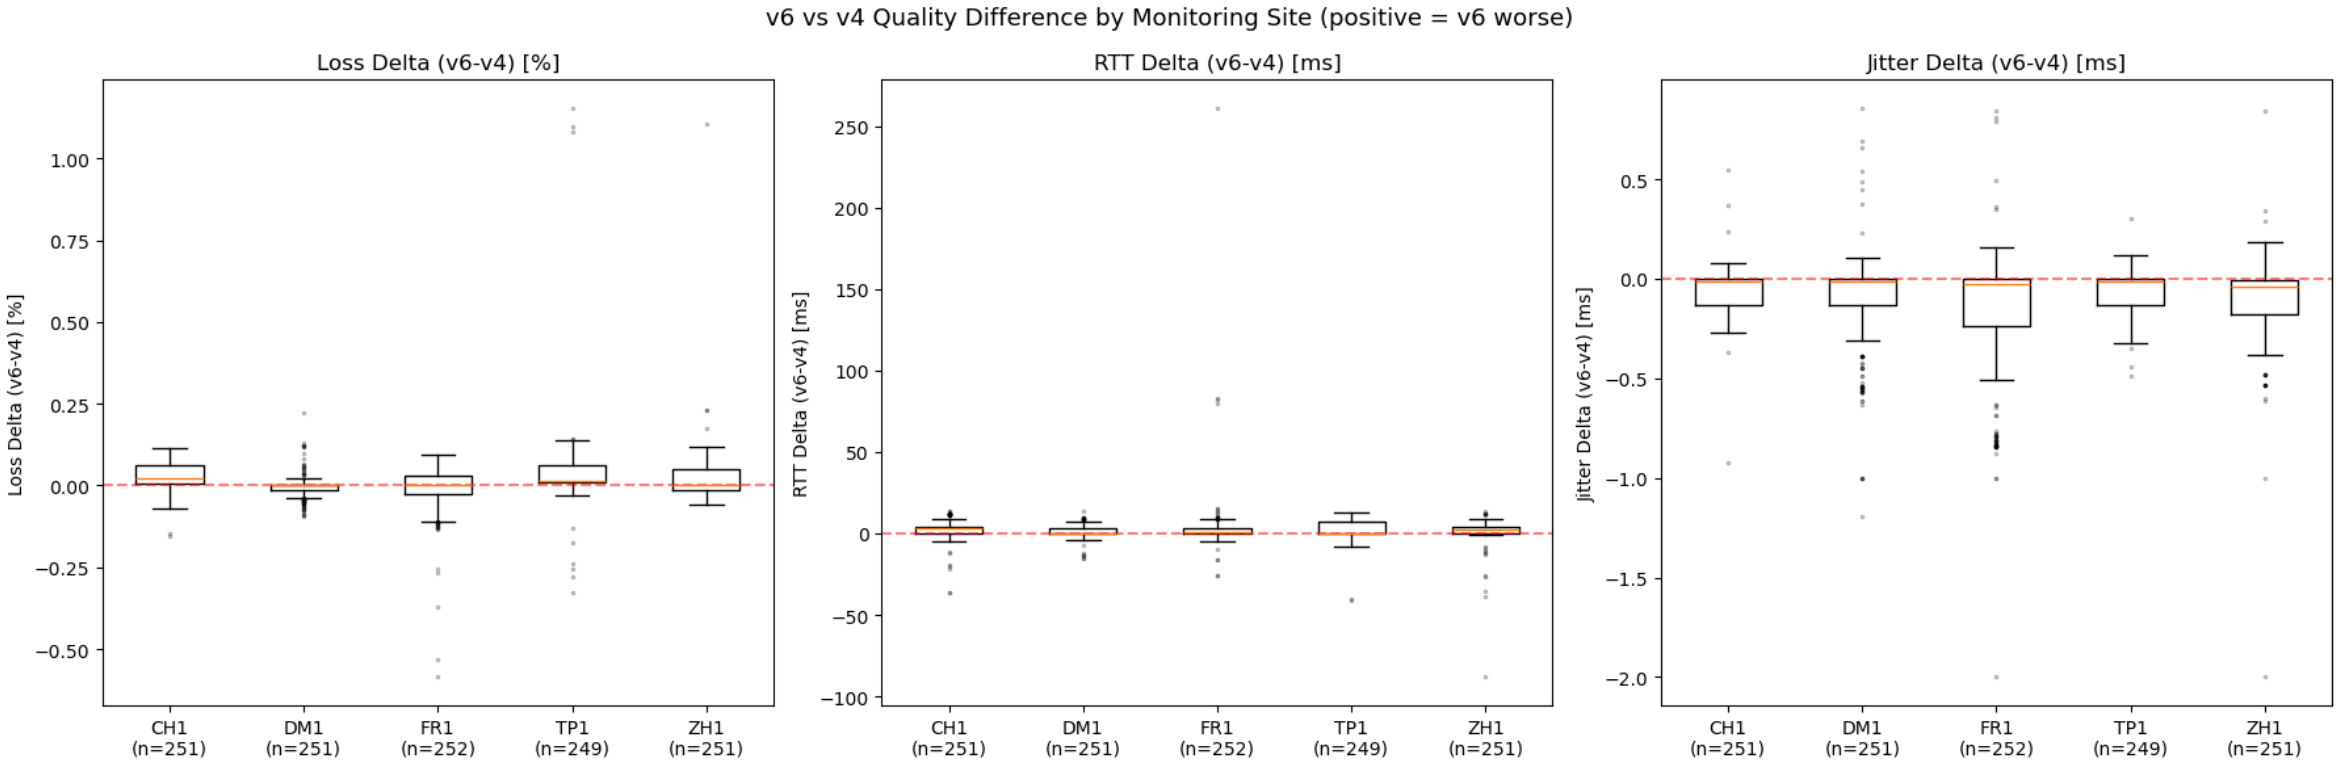

What goes wrong: IPv4

- Worse buffer-bloat — v4 show higher latency-under-load than their v6 twins.
- Stronger daily peak hours congestion pattern on v4.
- In-transit fragmentation artifacts — mid-path fragmentation that v6 simply forbids.
- Many paths are now faster on v6 — lower median RTT than the v4 equivalent.
- Generally worse jitter on v4 than on v6.
- *Address scarcity breeds broken design — stretched L2, NAT/IP pooling → larger portions of the network experience outage at once.*

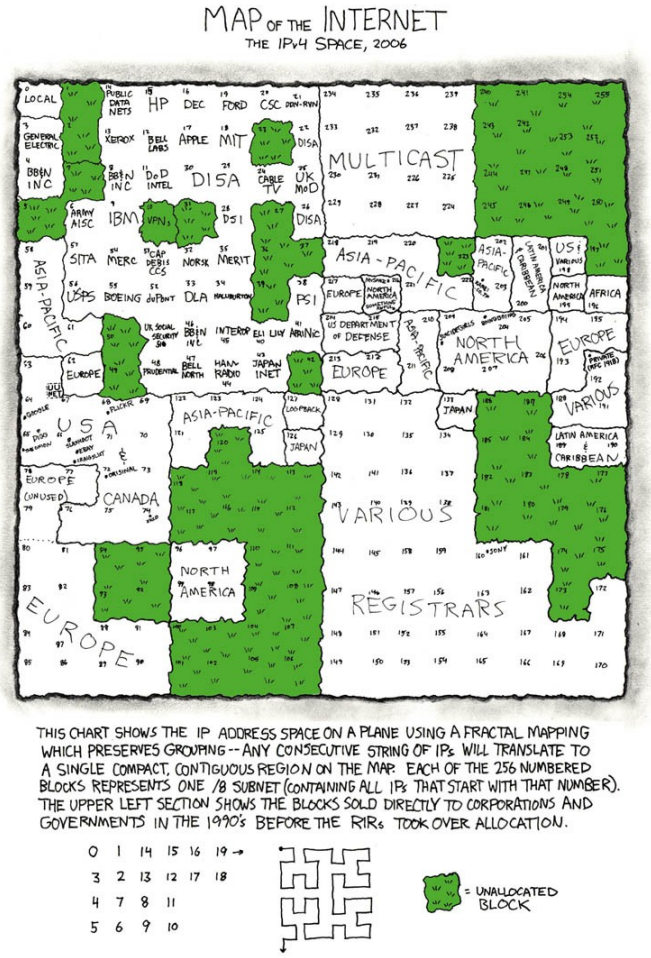

xkcd.com/195 "Map of the Internet — the IPv4 space" · CC BY-NC 2.5
"For the IPv6 map, just imagine the XP desktop."

Thank you!

Questions?