Can You Hear Me?

A New Approach to DNS Resolver Measurement

Maud Durand, Christian Huitema

OARC 47, Vancouver, BC
9 November 2026



Study Background – Same Data as Geoff

- This study is made in partnership with APNIC Labs
  - Based on APNIC Google Ads measurement system
    Thanks to Geoff Huston, Joao Damas & Christian Huitema
- **Initial goal: understand how ISP-based and PDNS-based resolution was used --in practice-- by ISP**
- Data collected every day since 4 November 2025 worldwide
  - **7,461,820,331 usable Unique IDs** (UIDs) or measurement points
  - Average about 900,000,000 measurement points per month
- This is still work in progress, more findings (and ICANN OCTO papers) to come!

Methodology Difference From Geoff's Presentation

- **We look at the data per ISP, no country level aggregation.**
  - **Note: some ISPs span multiple ASNs.**
- **We separated queries into DNS fluxes (a.k.a. DNS providers) based on the IP address of the resolver APNIC Labs is seeing:**
  - **ISP (IPv4 & IPv6)**
  - **Various PDNSs**
  - **"Others same CC", "Others"**
- We created time slices to count the duplicates per **RR type***:
  - **Initial query**: 0ms
  - **Immediate duplication**: <10ms, <30ms, <100ms
  - **Retransmit on timeouts**: <300ms, <1s, <3s, <10s, <30s
  - **Zombies**: > 30s

*we measure A, AAAA & HTTPS separately.

The Surprise: Very High Rate of Duplicates

- If we tally all categories, ISP, PDNS, and Others, 50% of all queries are repeated within 100ms.
- Another 30% are repeated within the next 30 seconds.
- These numbers are extremely high and totally unexpected.

Finding: the “Crossover Duplicate” Rate

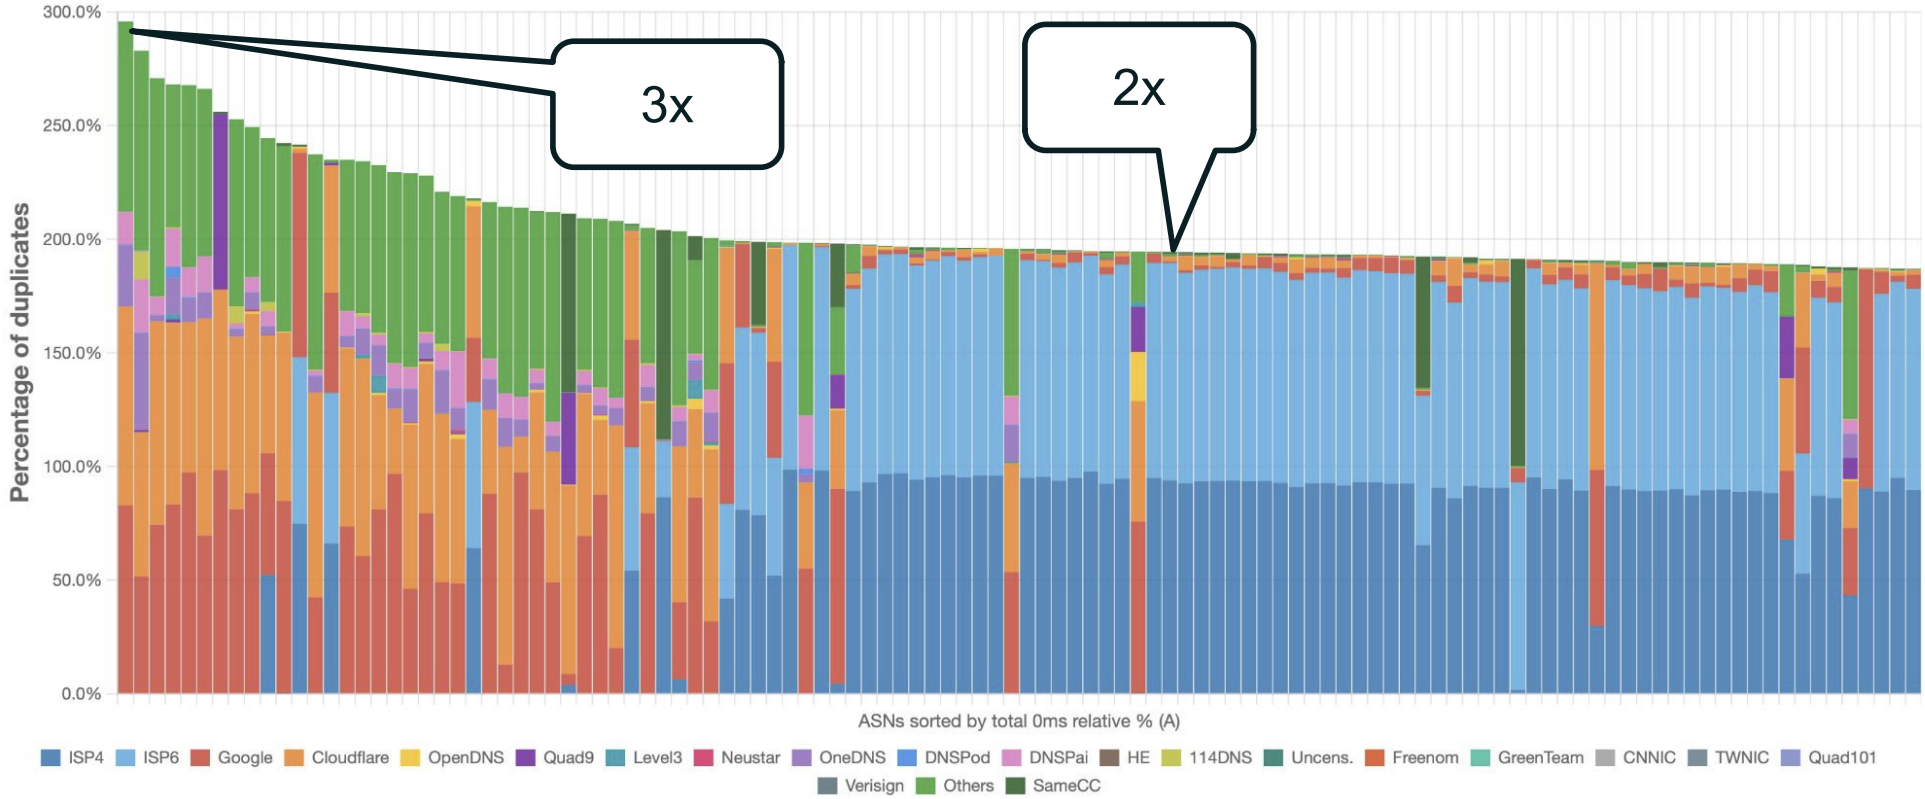

- Crossover: same query (name, RR) arriving through different paths:
  - Different DNS Providers, or ISP IPv6 + IPv4
  - **We count queries only once per flux**
  - Graph shows the “top” 3.3% of the distribution
    62% of ISPs have < 15% crossover.

Finding: the ”Global Duplicate Rate”

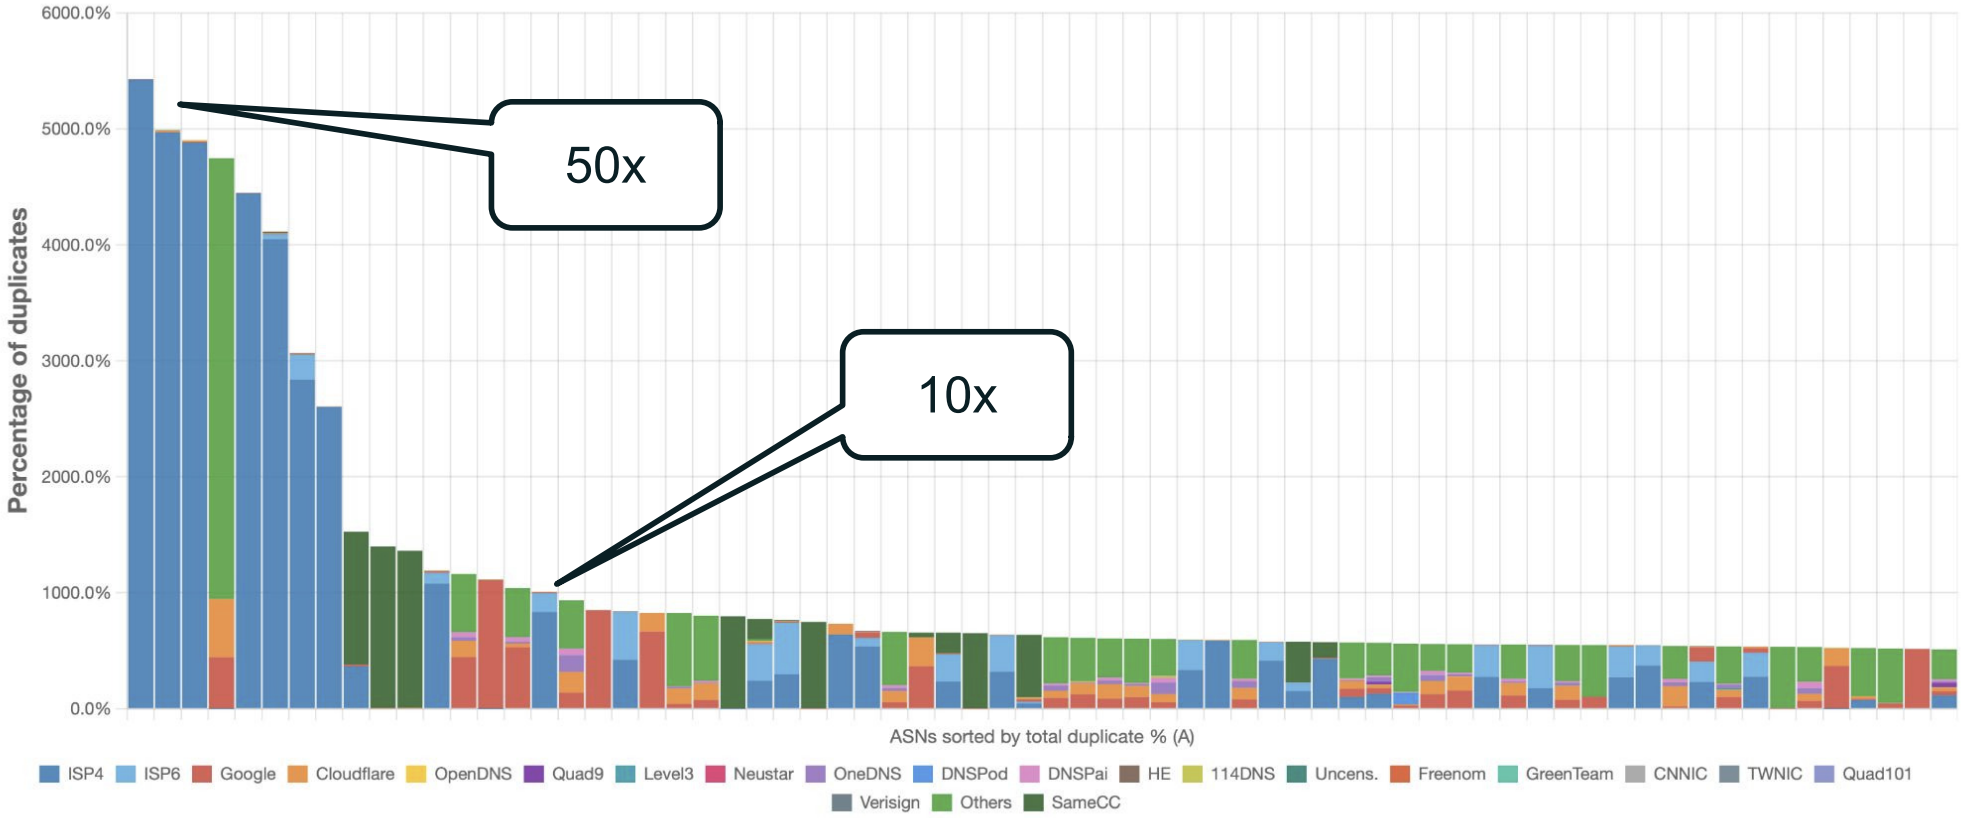

- **We count queries all queries in all fluxes**

- Graph shows the top 1.3% of the distribution
  About 58% of ISPs show less than 22% duplicates

Finding: Duplicate Rate Varies by Provider Type

- Different duplication rates for different DNS providers.
- Different DNS providers have different sizes, thus different impact.
- **80.9% of all duplicates come from ISP-originated queries**

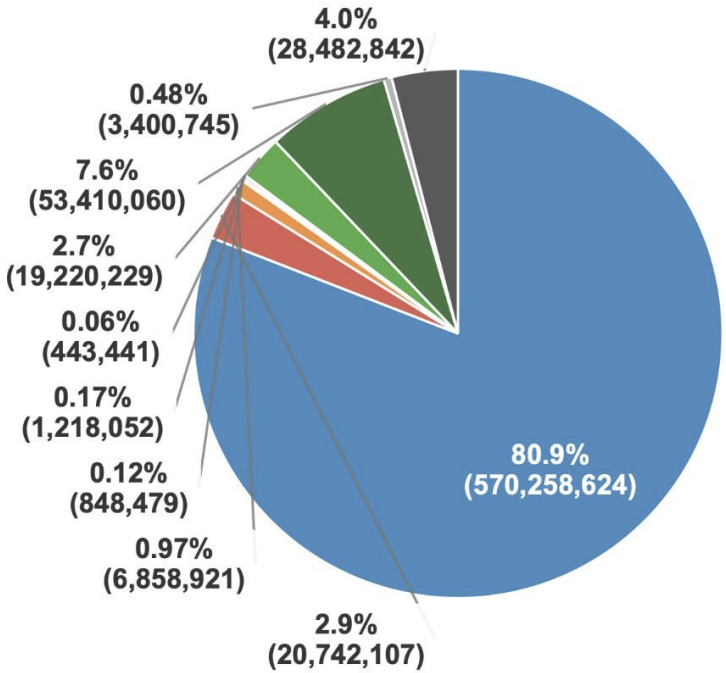

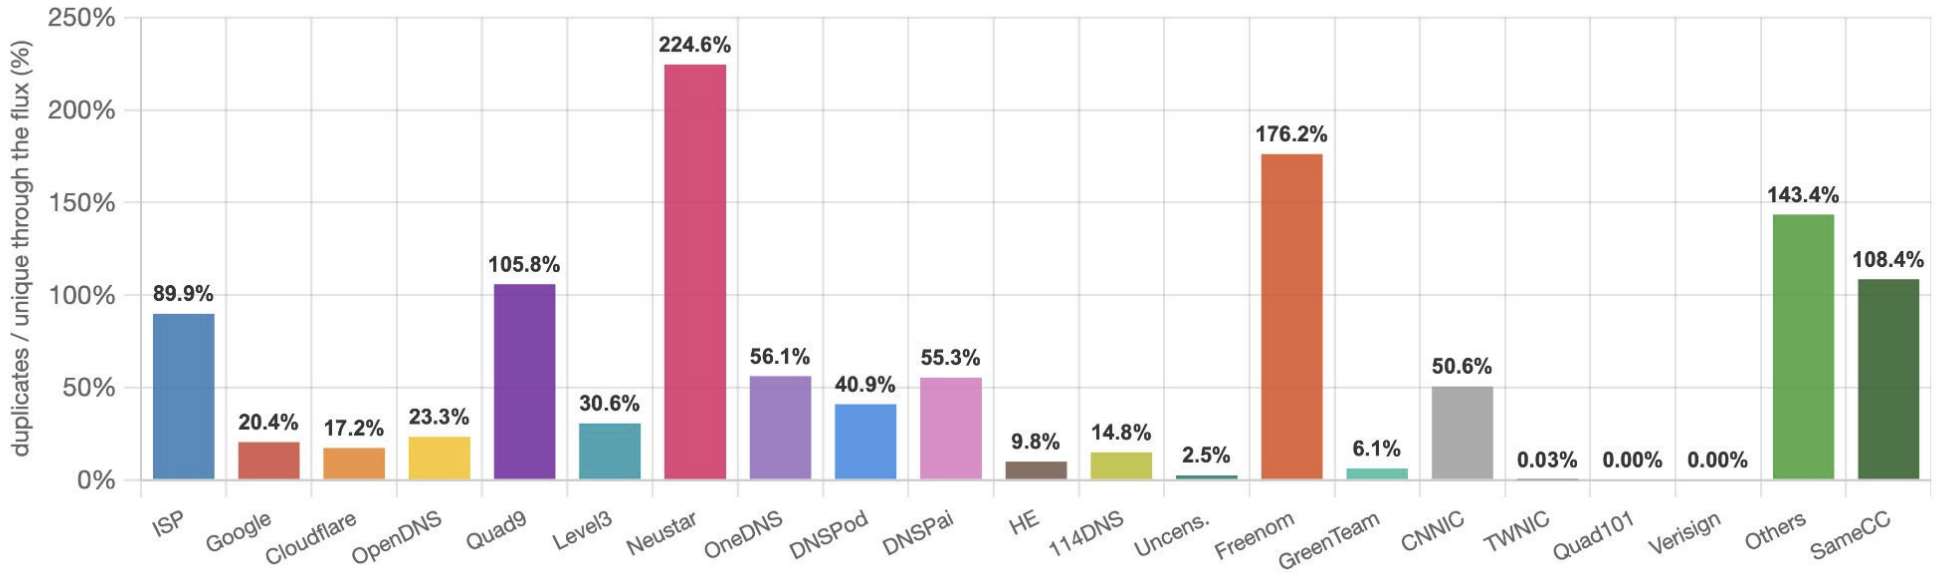

ISP Patterns

Examples

Example 1: Baseline, Simple Pattern

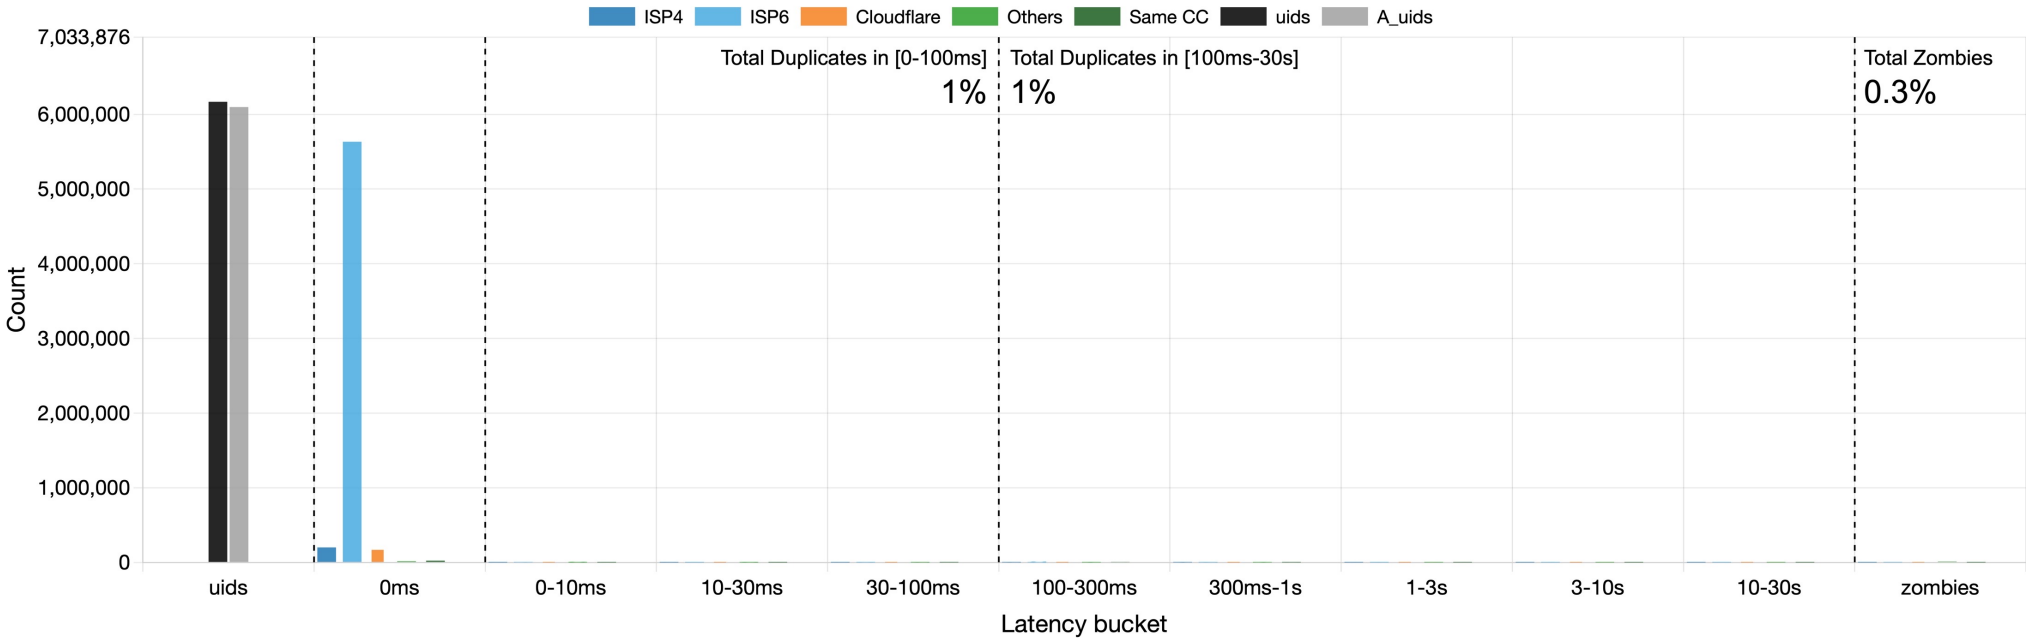

In this example:

- 92.5% of the first queries are received through an IPv6 address allocated to the ISP,
- 3.35% through an IPv4 address of the ISP
- The remaining through Cloudflare and other public DNS recursive servers.
- We don’t see much duplication after the first queries.

- **The total of duplicates is just 1.4% of traffic.**

Example 2: Belt and Suspenders

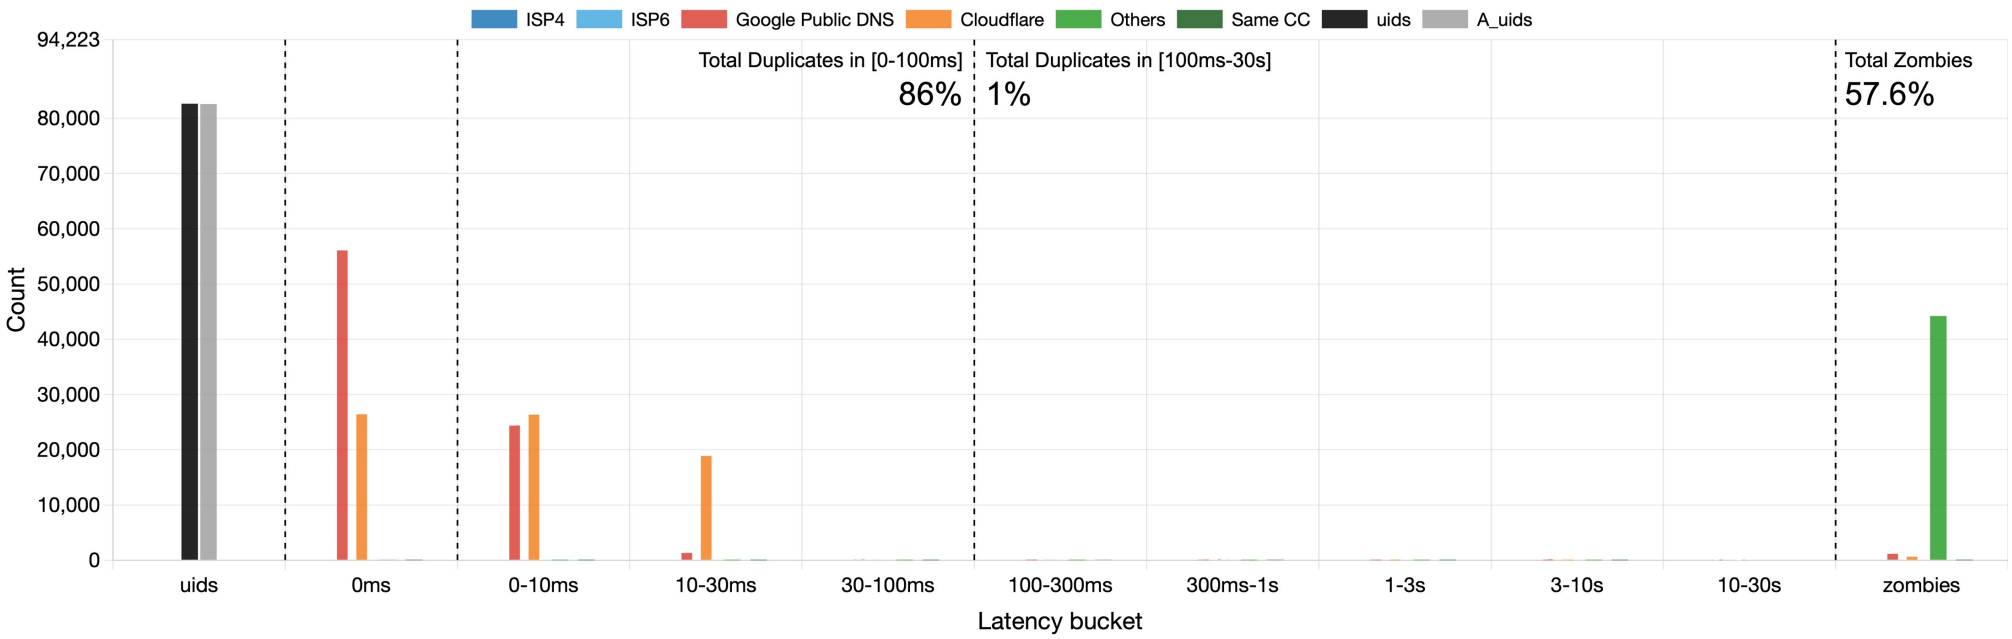

- Most queries are received from both Google Public DNS (99.2%) and Cloudflare (87.6%).
  - Solid bars in graph: first arrival from that DNS Service Provider,
  - Graph shows duplicates very shortly after 1st arrival. (network delay?)
- **High ratios point to a decision in the ISP.**
- Belt and suspenders: whatever response arrives first wins?

* With some delays…

Example 3: Happy Eyeball Pattern

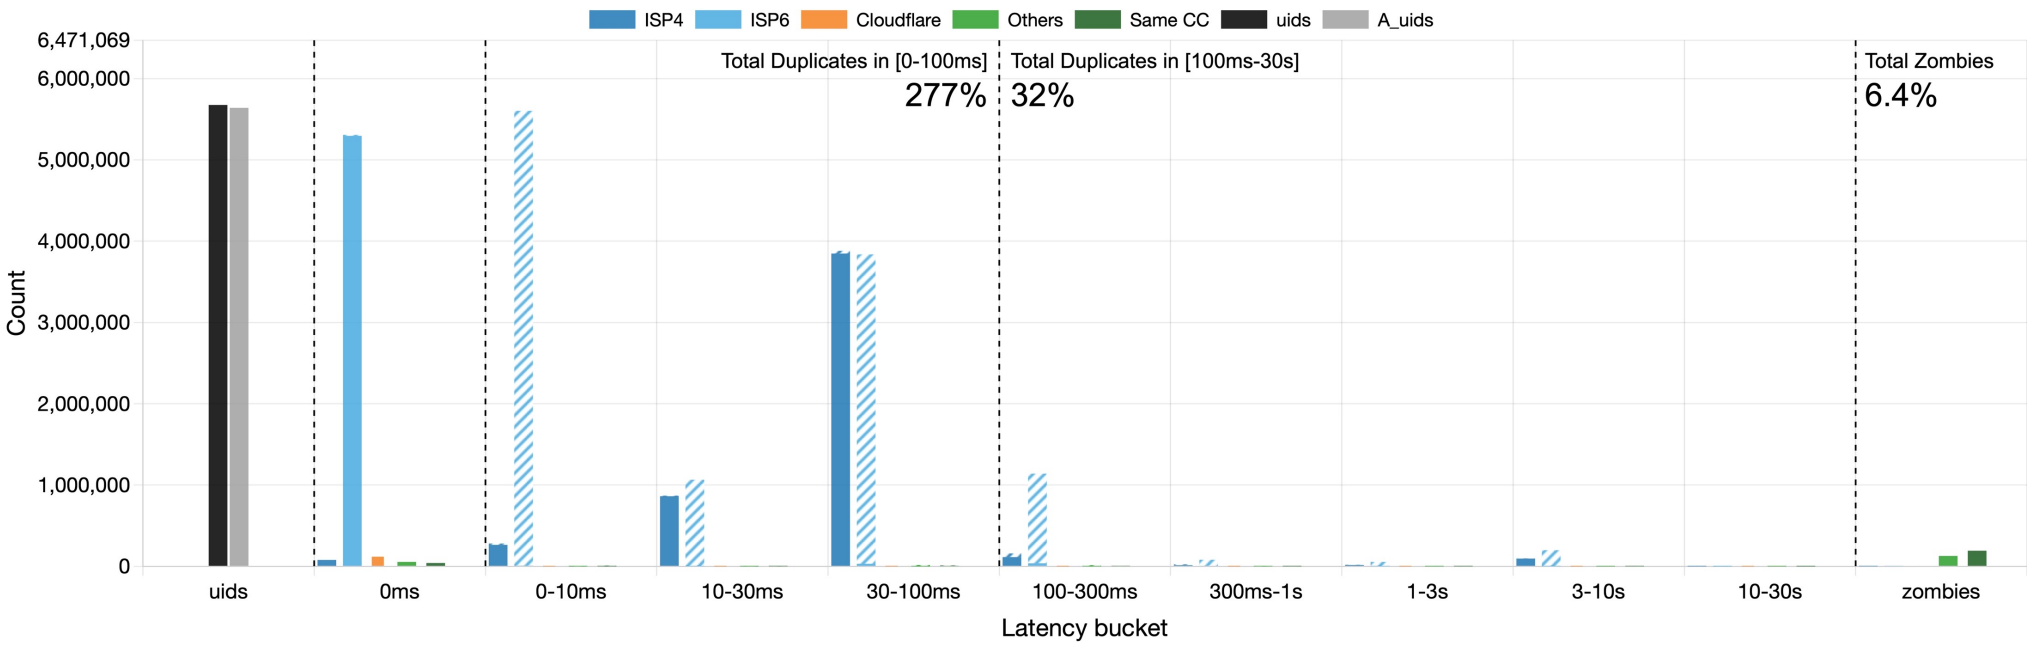

- **Most queries duplicated over IPv4 (93.9%) and IPv6 (95.6%)**
- Lots of repeated IPv6 queries
  - Hashed lines in graph show duplicate arrival from same source
  - 194.5% duplicate overall
- Named “happy eyeball” pattern by analogy with RFC8305, but there may be another explanation.

Example 4: Heavy retransmit

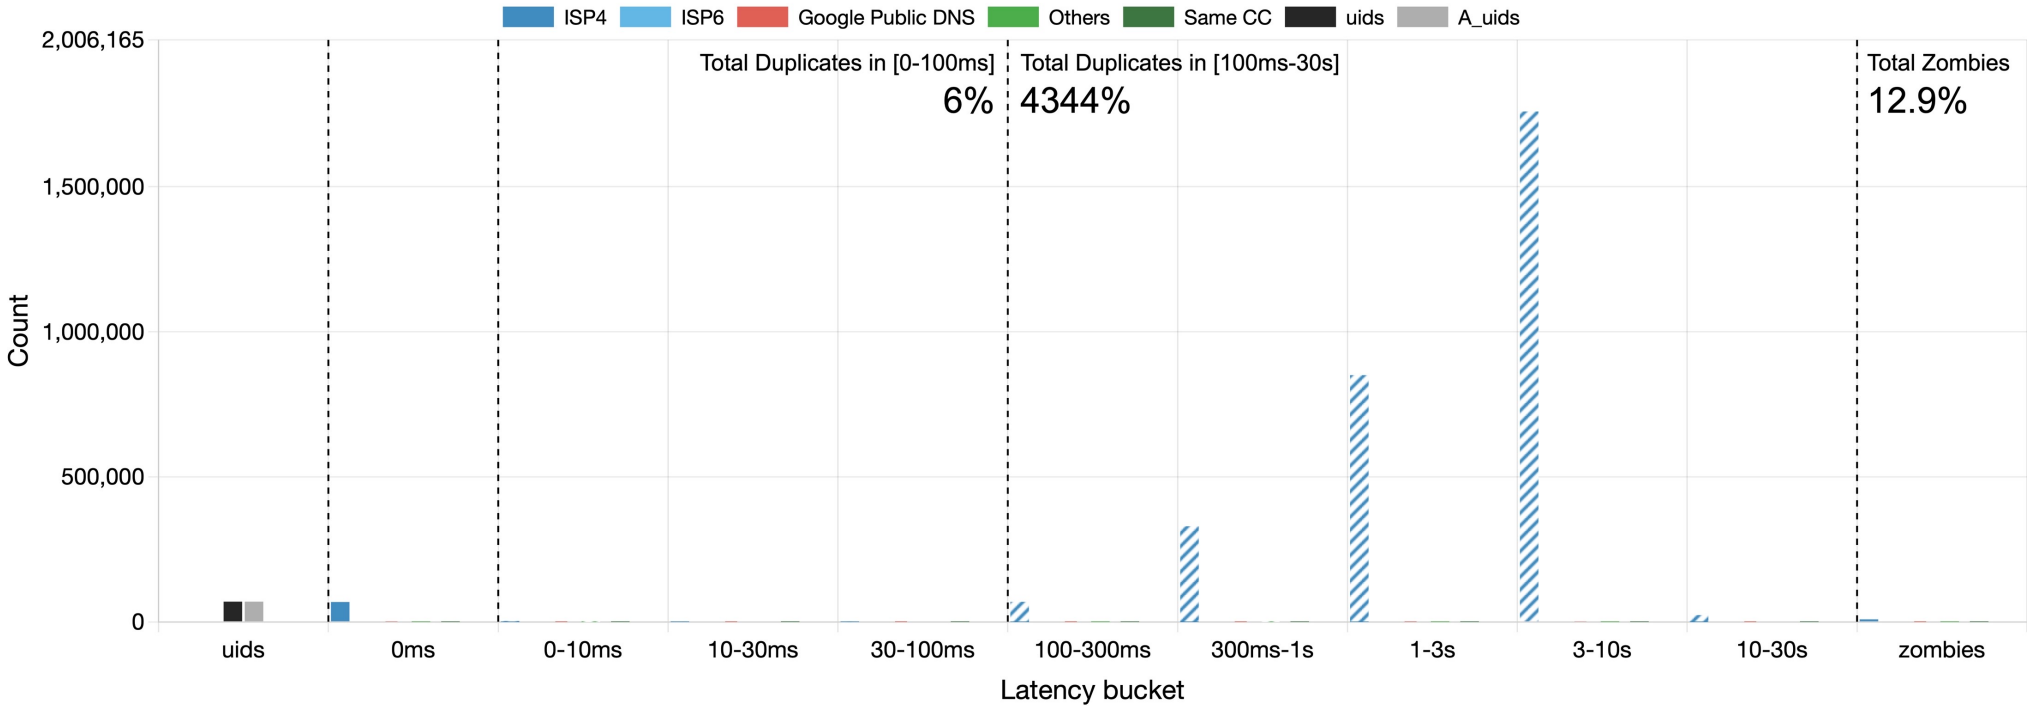

- The 100ms-300ms, 300ms-1s, 1s-3s show duplicate arrival rates proportional to length of interval
  - Consistent with repeat every 200ms for 5s.
- Clearly not the exponential backoff suggested in RFC 8085.
- Each query is repeated more than 25 times, clearly some kind of bug.

ISP Patterns

Analysis

Analysis: How Prevalent Are Each Patterns?

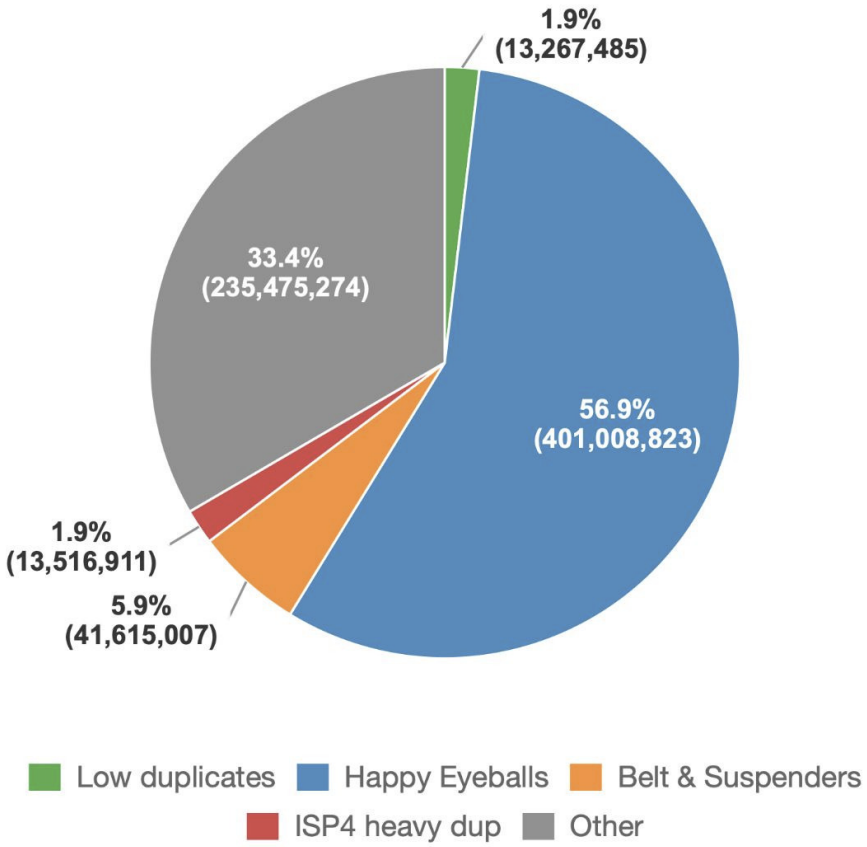

	%dup	%ISP	%Uniq.
Low Duplicates	**1.8%**	24.7%	31.5%
Happy Eyeballs	**56.9%**	5.7%	23.5%
Belt & Suspenders	**5.9%**	10.9%	3.7%
Heavy Dup.	**1.9%**	0.2%	0.04%
Others	**33.4%**	58.6%	41.2%
Total	**100%**	100%	100%

Analysis: Why So Many Happy Eyeballs Duplicates?

- Happy Eyeball RFC8305 could introduce a x2 duplication rate
  - We see duplication rates > 3 and sometimes > 4 across IPv4 and IPv6
  - We spoke to ISPs that pointed us to ‘policy enforcing stateful firewalls’
- The firewall hypothesis
  - Client performs DNS query, client loads web page
  - Firewall intercepts web page, extracts domain name
  - Firewall checks request for compliance with ISP policies.
- That leads to x2 DNS queries if the firewall does not use the same DNS server infrastructure as the ISP subscribers.
  This explanation, compounded by actual happy eyeballs RFC8305, could explain the x3 or x4 observed duplication.
- We need to perform further analysis and engage with large service providers.

Open Questions

- **How much of this is an ISP choice vs a software bug?**
  - **Belt & Suspenders**: looks like an ISP decision
  - **Happy Eyeball**: to be investigated further
  - **High level of duplicates**: Appear in a number of ISPs in different geographies, potentially a bug in a widely used piece of software

- **Zombies**
  - There are also discernable patterns
  - It looks like zombies are often replayed in different networks than the ISP
  - Sometimes the replay rate is high
  - For further studies

Engage with ICANN

Thank You and Questions

Visit us at **icann.org**
Email: email


@icann


facebook.com/icannorg


youtube.com/icannnews


flickr.com/icann


linkedin.com/company/icann


soundcloud.com/icann


instagram.com/icannorg