A Look at Traffic to Authoritative DNS Servers of a Large Enterprise

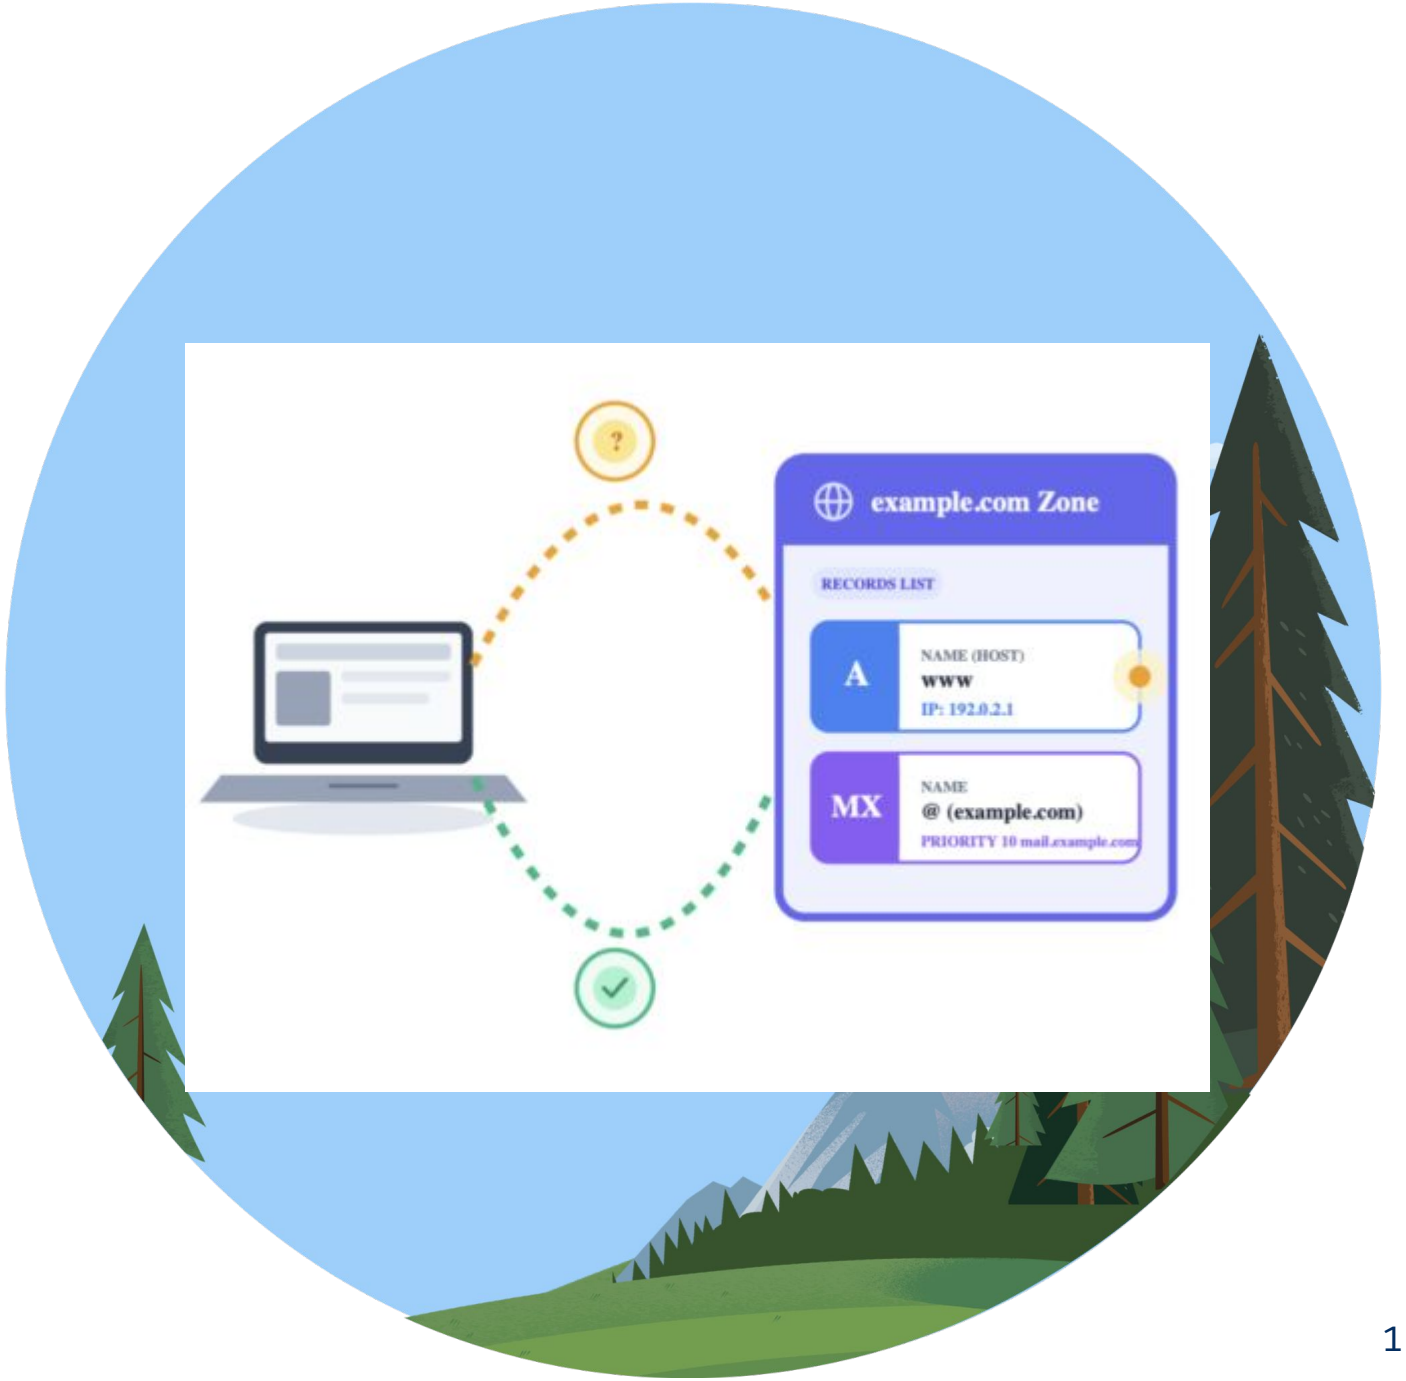

Pallavi Aras & Shumon Huque

DNS-OARC 46 Workshop
May 16th 2026
Edinburgh, Scotland

Presentation Goals

Actionable Insights

- NX domain analysis
- Geo distribution
- Top 10 carriers

Trends

- Top 10 carriers
- Response code distributions
- UDP/TCP truncation

Characterization

- Traffic analysis
- DNSSEC

Architecture

Comprehensive data collection framework review.

Takeaways

Analysis utility and provider recommendations.

Data Foundation & Analysis

Database & Source

Uses raw query data captured and stored by DNS provider, over an extended timeframe.

Presentation Scope

Data analysis is performed using Anthropic Claude to get insights and data correlation.

Full **data collection architecture** will be detailed in the concluding section.

Actionable Insights

NX Domain Analysis

NX Domain Analysis - Goal

Objectives

- Determine which zone consistently receives the highest volume of NXDomain queries
- Identify the underlying causes and indicators of these query patterns

Primary Finding

Email Infrastructure Domain

Recorded the highest volume of NXDOMAIN results across the analyzed zones

NX Domain data analysis

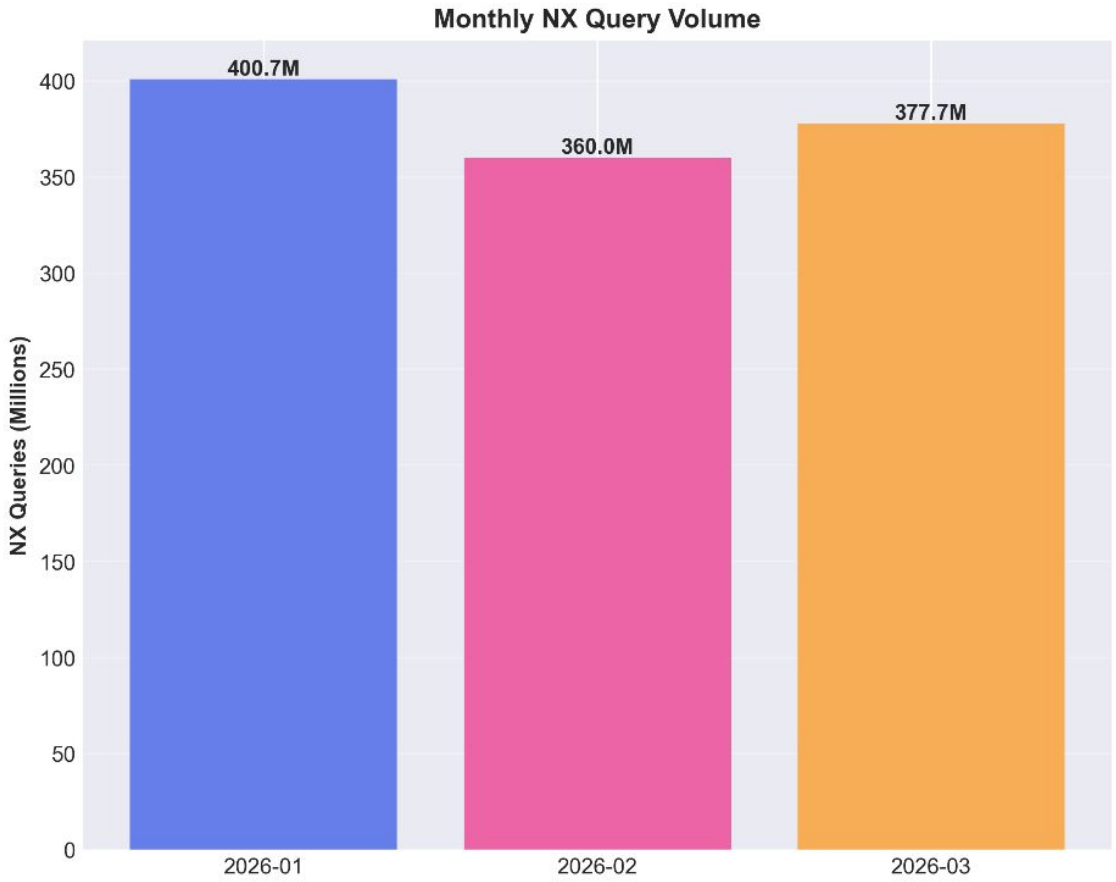

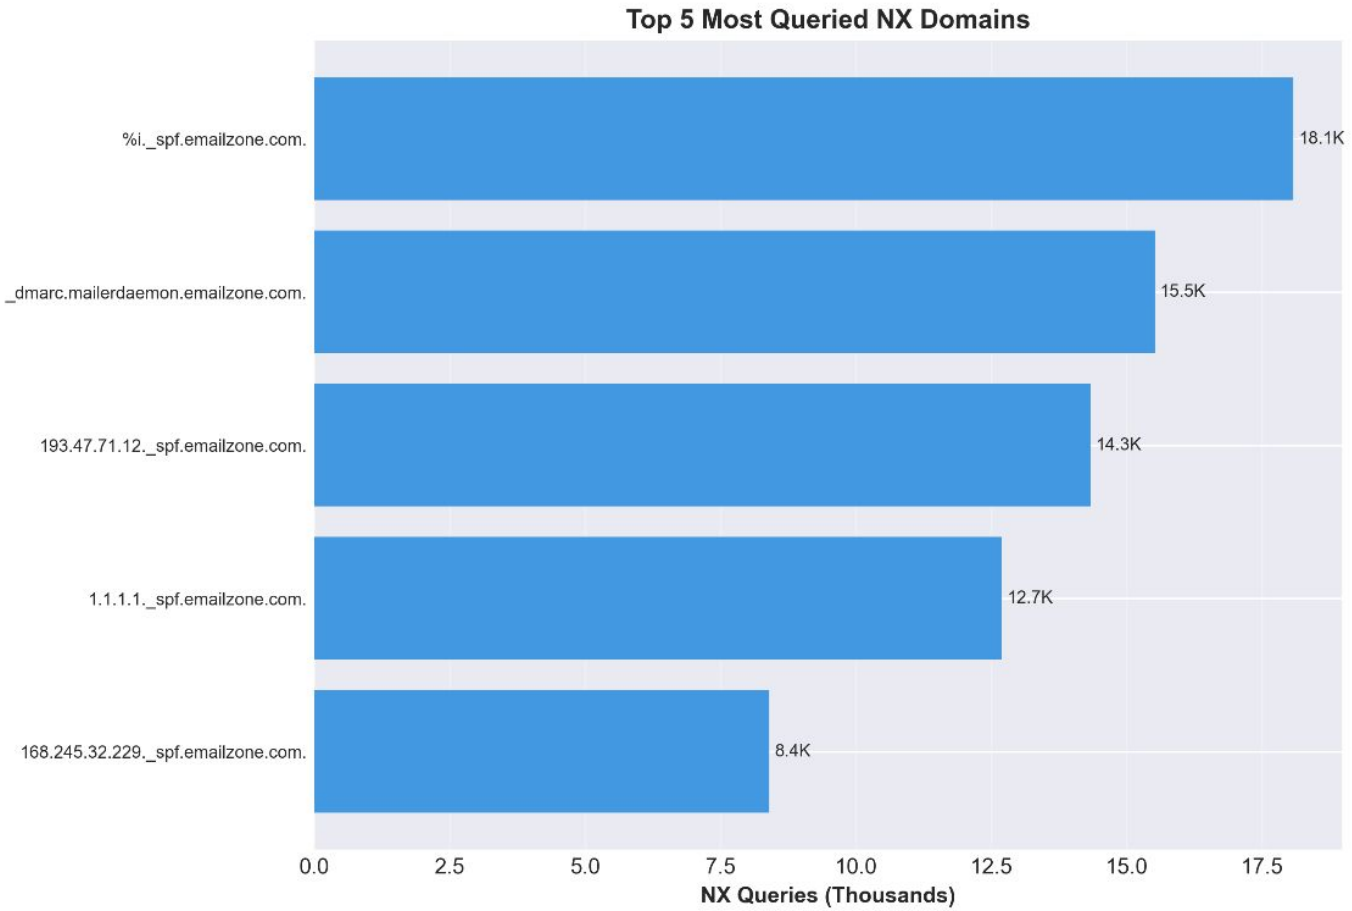

NX Domain Data Analysis: <emailzone.com>

Pattern Analysis

1. %i._spf.<emailzone.com>

- **32,516** variable pattern queries
- "exists" clause in SPF record

2. DMARC Misconfigurations

`_dmarc.mailerdaemon<emailzone.com>`

Domain doesn't exist but is heavily queried.

Question: WHY?

IP-Based Dominance

IP-based queries dominate the traffic

778M queries for domains like:
`193.47.71.12._spf.<emailzone.com>`

68.4%
IP-based SPF Queries

Possible misconfiguration or malicious reconnaissance activity.

Actionable Next Steps

Work with **Email Infra** to correlate the data with misconfigurations and incorrect query patterns

Work with **Security team** to determine potential recon attacks.

Geographic Distribution of DNS Queries

Geographic Distribution of DNS Queries

UltraDNS Pods - for main critical zone

116.3B
Total DNS Queries Analyzed

43 / 12
Locations across Regions

3.75%
Overall NXDOMAIN Rate

Top Location: uswas1 (Washington DC) with 18.3B queries
Top Region: North America East with 25.8B queries
IPv6 Adoption: 38.7% in Europe-North (Highest) vs 2.2% in Asia-West (Lowest)

Geographic distribution of DNS queries across Vendor nodes

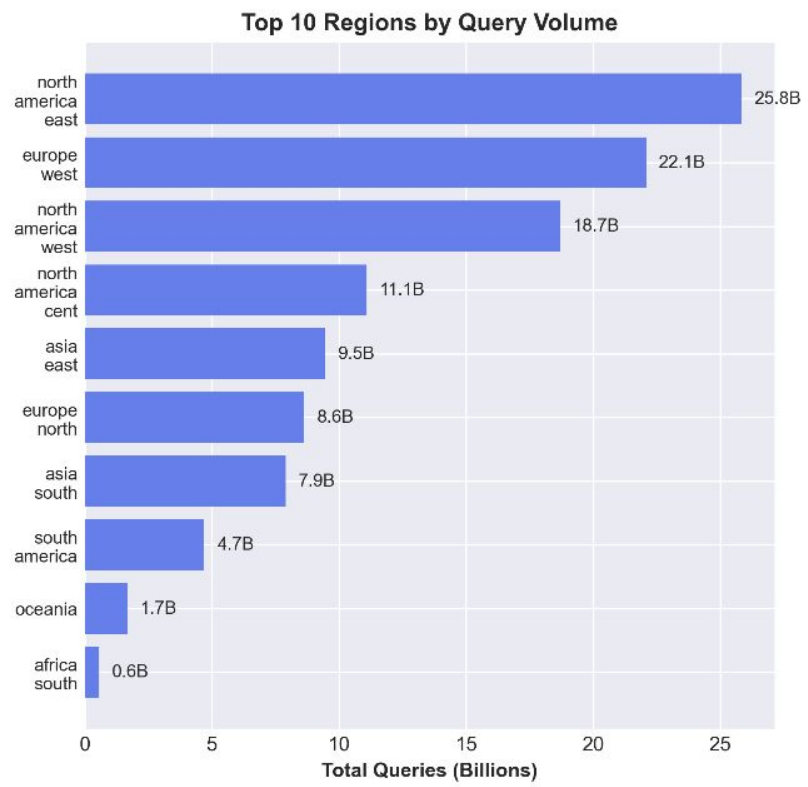

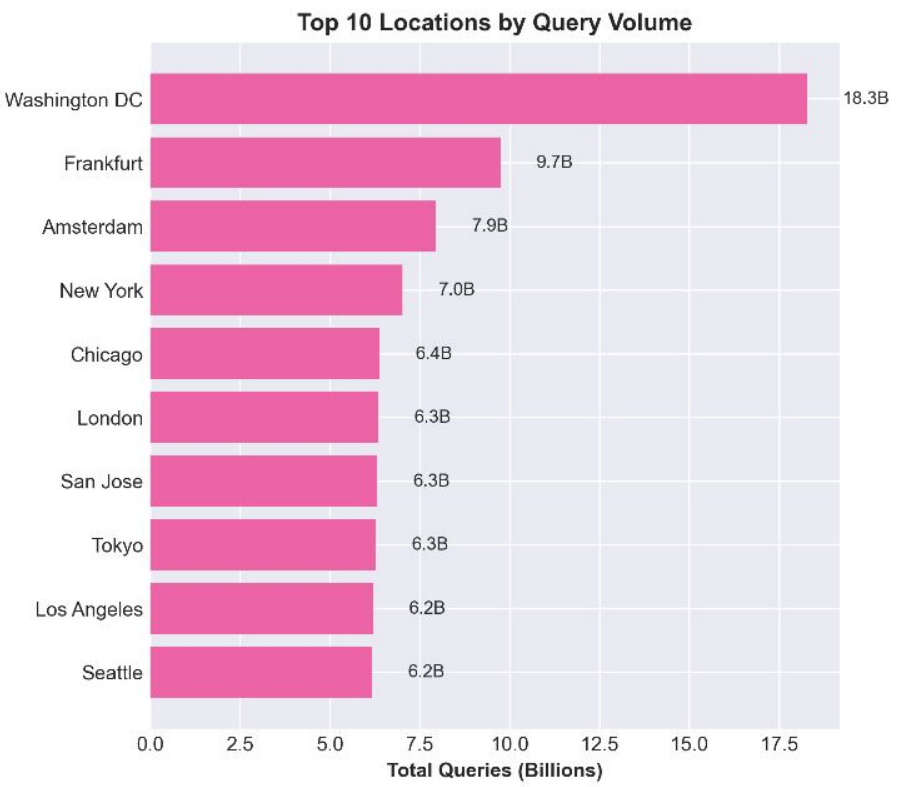

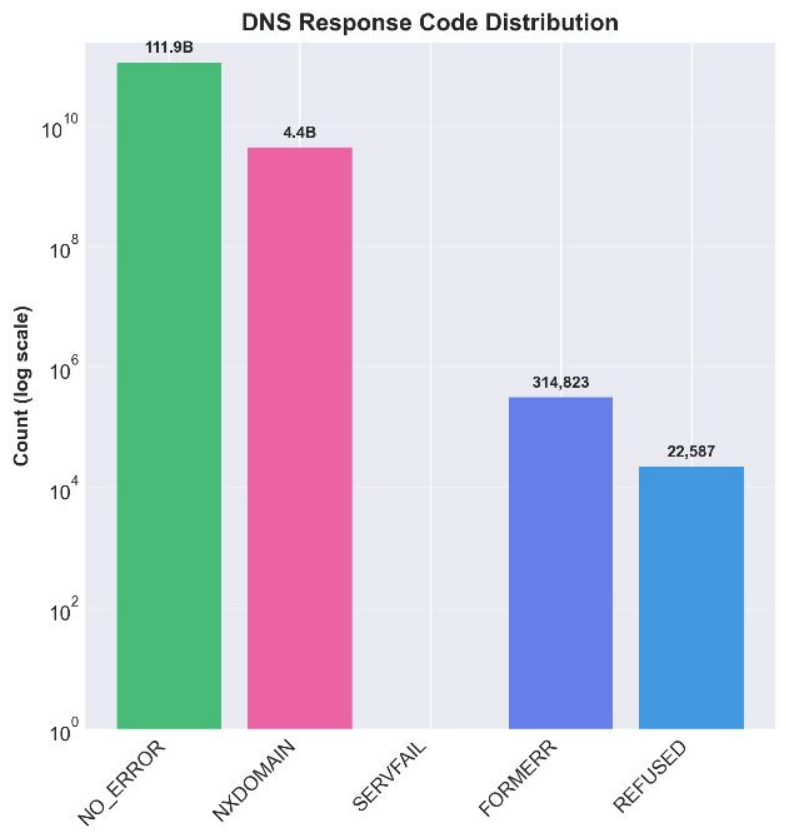

Geographic distribution of DNS queries across Vendor nodes

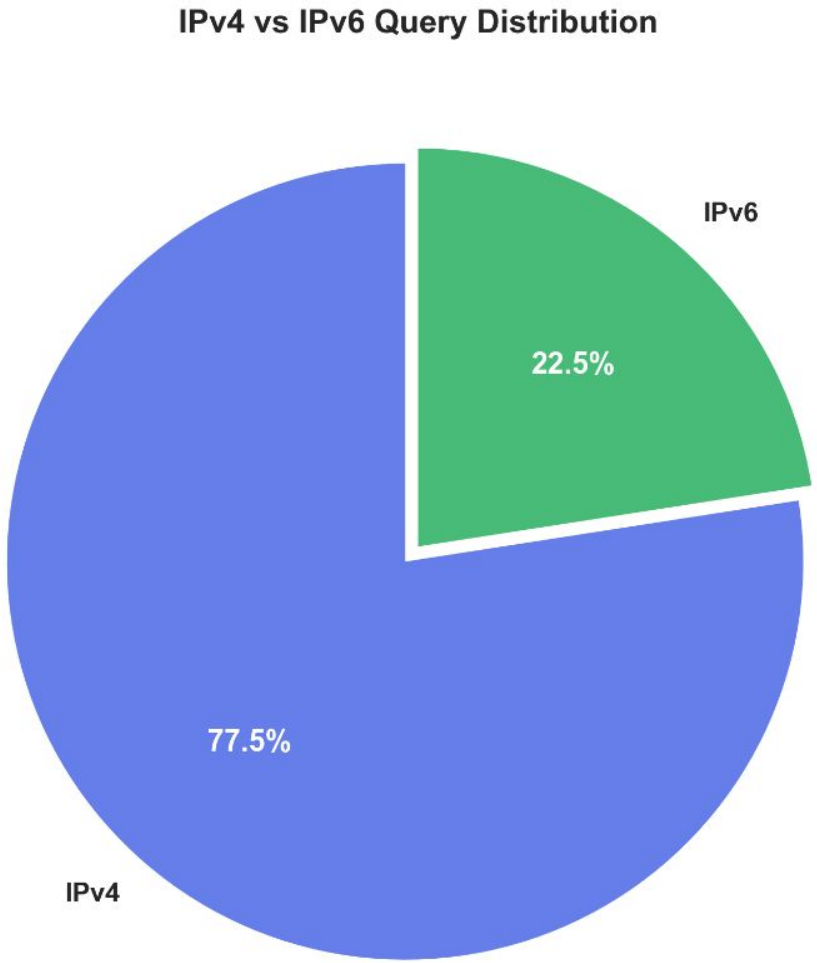

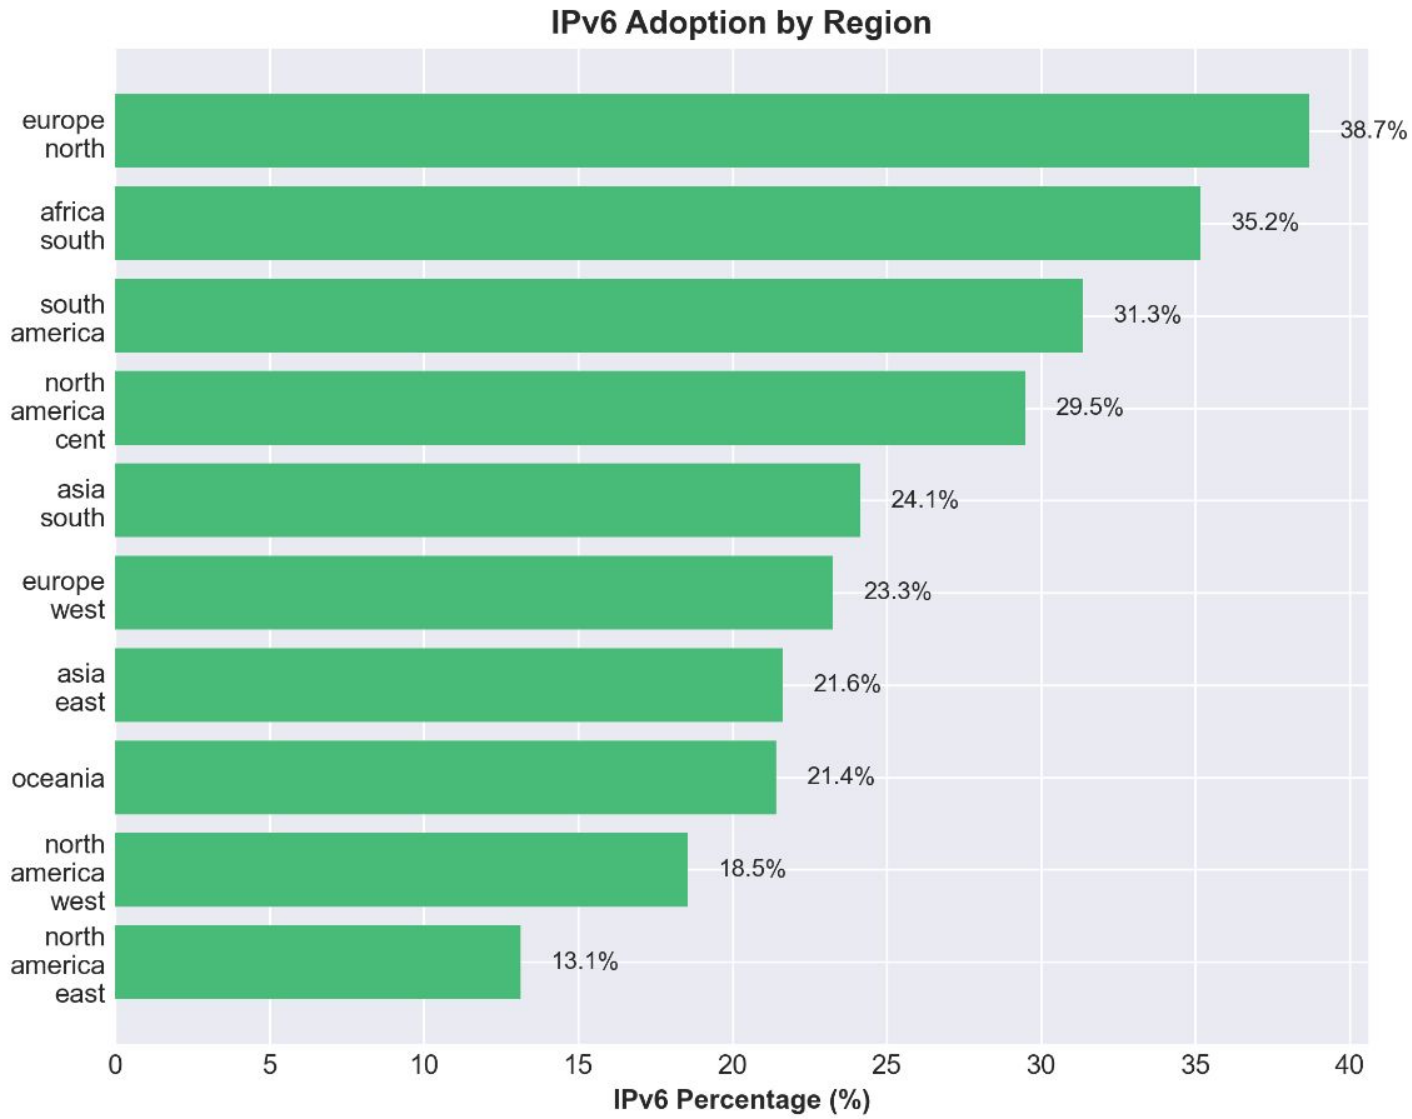

Actionable next steps

- Evaluate **capacity** in the most heavily used geo-locations.
- Global load balancers can be designed to **reduce latency**.
- Deployment of **IPV6 adoption** in EU can be prioritized.

Country-based Query Distribution for Zone [EDNS Client subnet(ECS) v/s ResolverIPs]

Country-based Query Distribution for Zone- ECS

90.1B [9 months]
Total DNS Queries Analyzed

5 Zones
Analyzed DNS Zones

34
Unique Countries

Most queried zone: livechatzone.example (30.6B queries, 34.0%)
Most active country: United States (32.6B queries, 36.2%)
Top 5 countries: United States, Singapore, Hong Kong, India, United Kingdom,Germany

Country-based Query Distribution for Zone-ECS

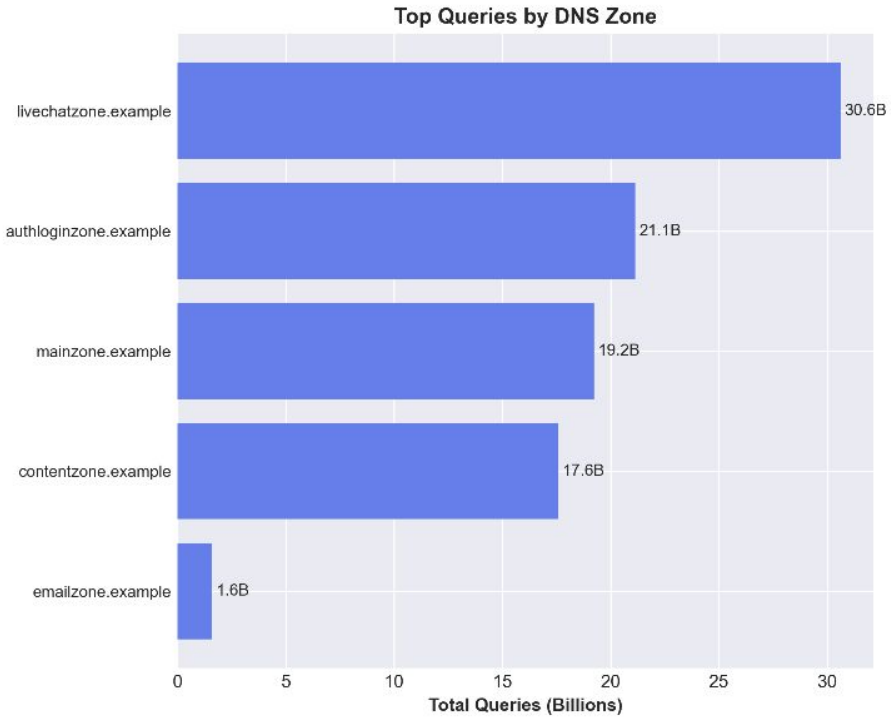

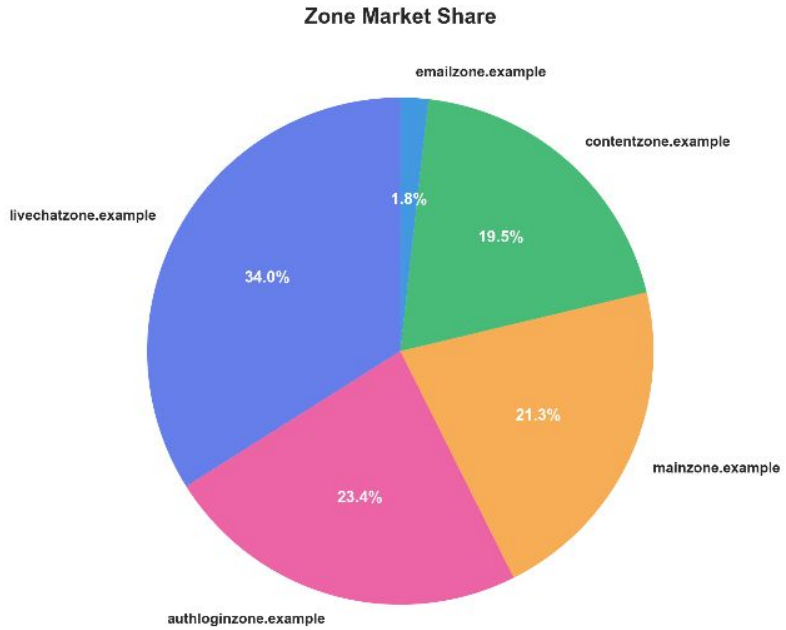

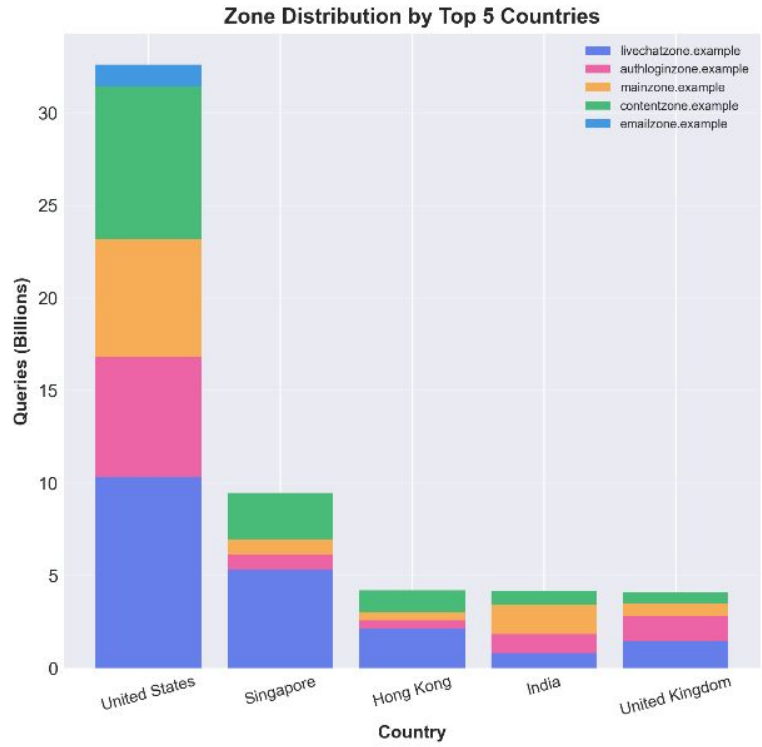

Country-based Query Distribution for Zone - ResolverIPs

1.26 trillion[9 months]
Total DNS Queries Analyzed

5 Zones
Analyzed DNS Zones

28
Unique Countries

Most queried zone: **authlogizone.example** (637.0B queries, 50.5%)
Most active country: United States (783.4B queries, 62.0%)
Top 5 countries: United States, Germany, Japan, India, United Kingdom

Country-based Query Distribution for Zone - ResolverIPs

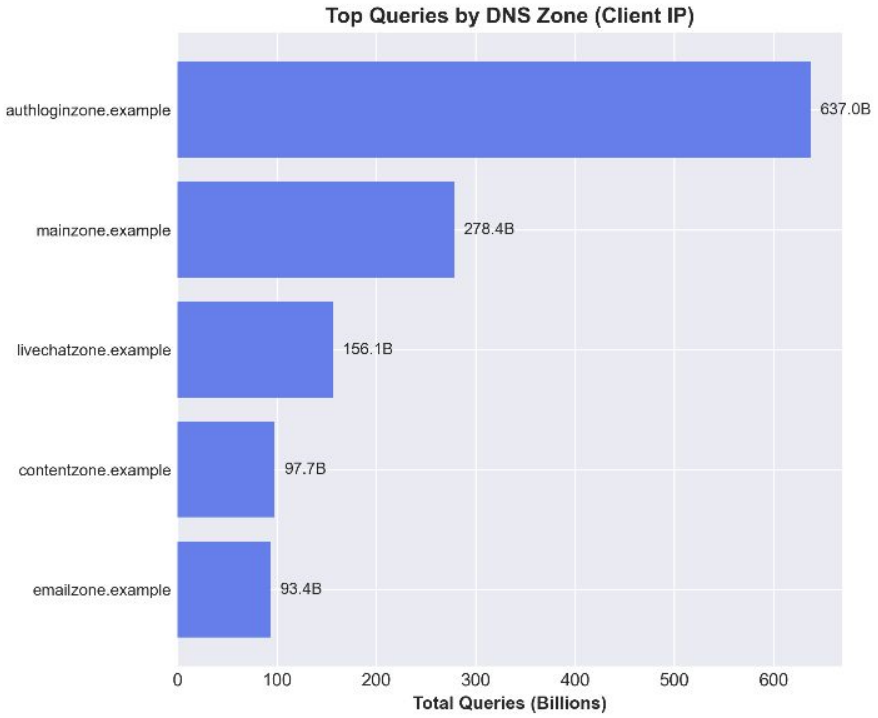

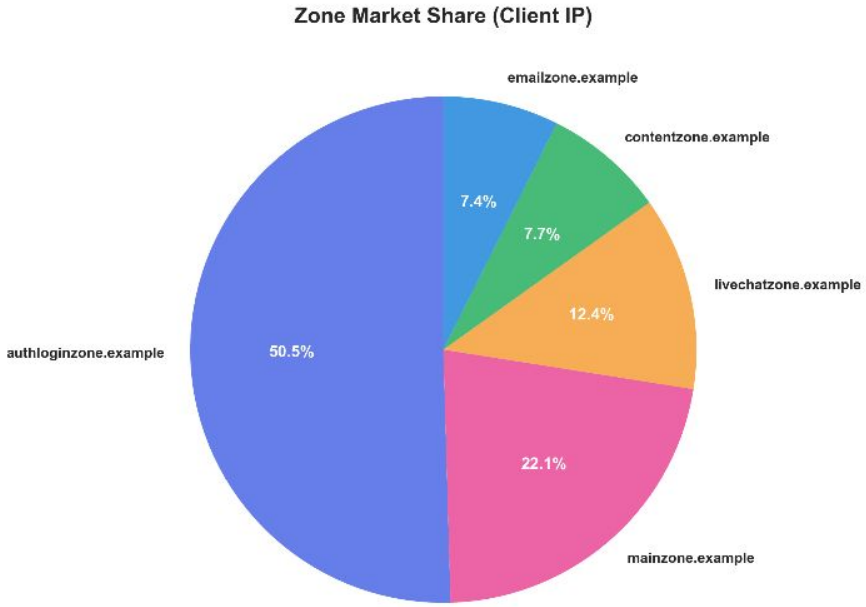

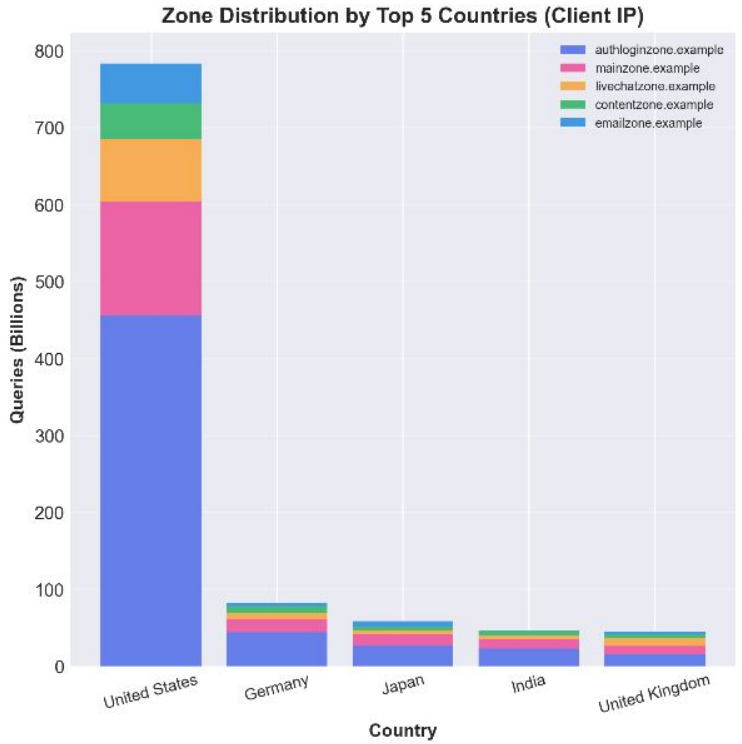

Actionable next steps

Analyze deeper why Zone popularity varies based on inclusion of ECS.

Application footprint improvement based on Zone popularity

TOP talking Clients
[EDNS Client subnet(ECS) v/s ResolverIPs]

TOP talking Clients - ECS

Total Queries (Last 3 Months)

19 Billion [Clients sending ECS option]

SCloud

9.7M queries/subnet

399M queries from 41 subnets.

Very concentrated traffic

UCloud

5.2M queries/subnet

2B queries from 394 subnets.

Highly concentrated

Comcast

179 queries/subnet

357M queries from 1.99M subnets.

Very distributed traffic

What this tells us:

High queries/subnet (SCloud, UCloud): A few IP ranges generating massive volumes - likely data centers, APIs, or automated systems.

Low queries/subnet (AT&T, Comcast): Many IP ranges with low individual volume - typical of residential ISPs.

TOP talking Clients - ECS

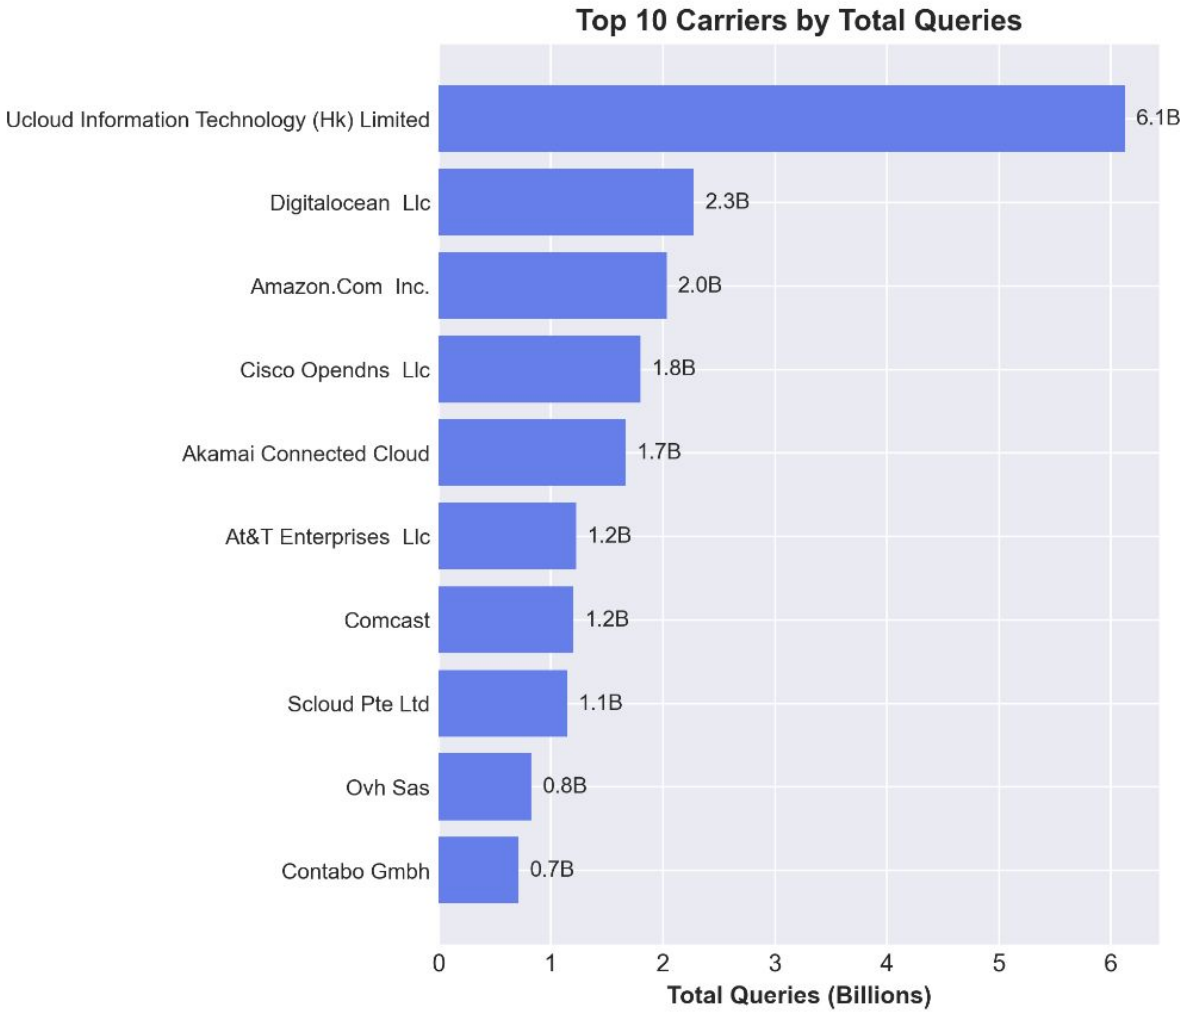

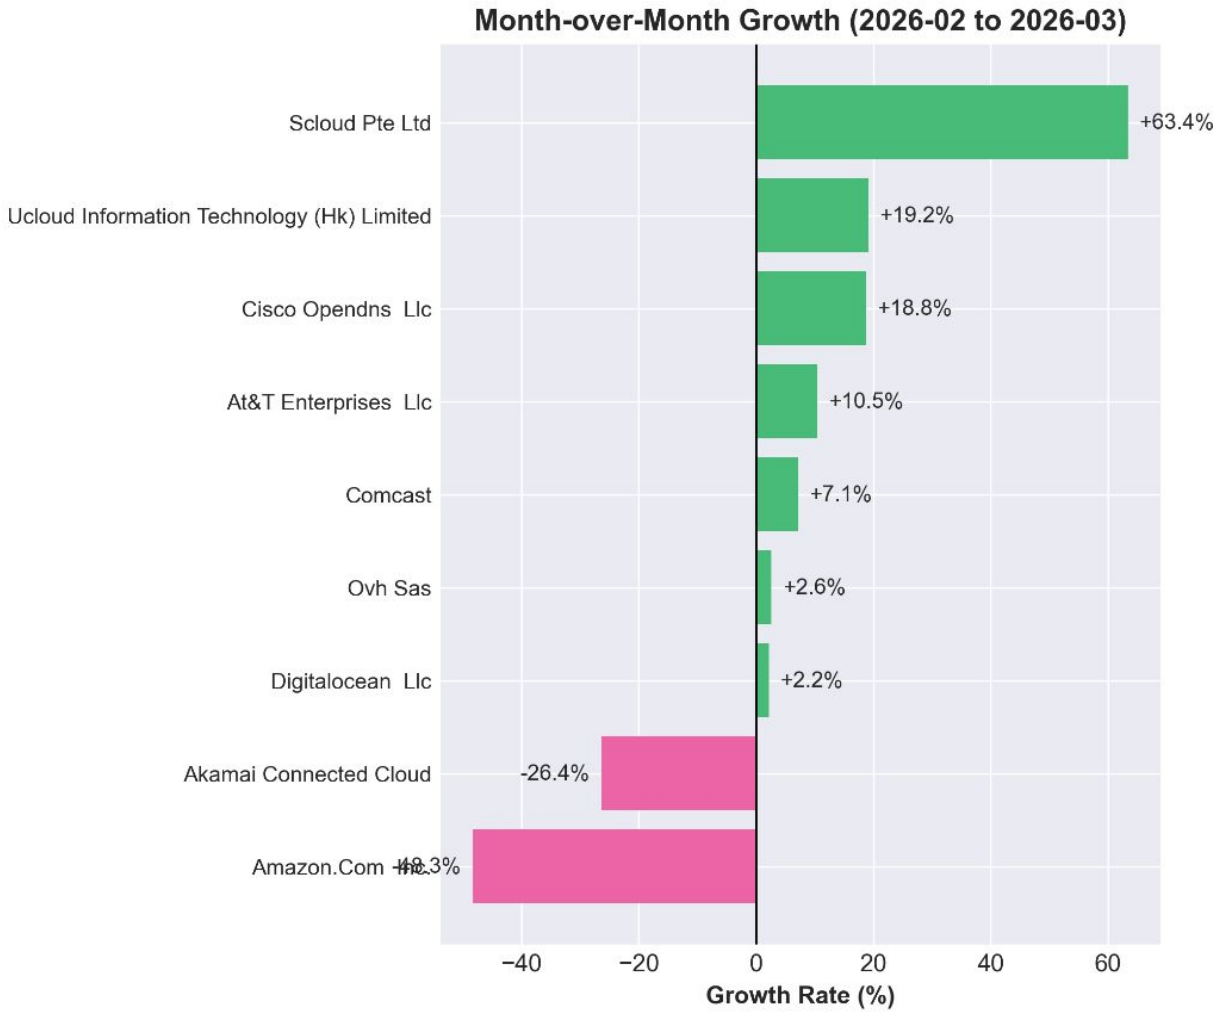

TOP talking Clients - ResolverIPs

Total Queries (Last 3 Months)

470.6 Billion [ClientIP analysis]

Salesforce

8.7M queries/IP

Highest concentration - internal infrastructure.

Cisco OpenDNS

2.4M queries/IP

DNS resolver service.

Fastly

1.7M queries/IP

CDN infrastructure.

Key Client IP Insights:

Scale: 727,407 total distinct client IPs in March, with an average of 186,662 queries per IP.

Cloud/CDN Dominance: 93.6% of all queries originate from major cloud and CDN providers.

Concentration: High per-IP query volume observed from company network infrastructure.

TOP talking Clients - ClientIP

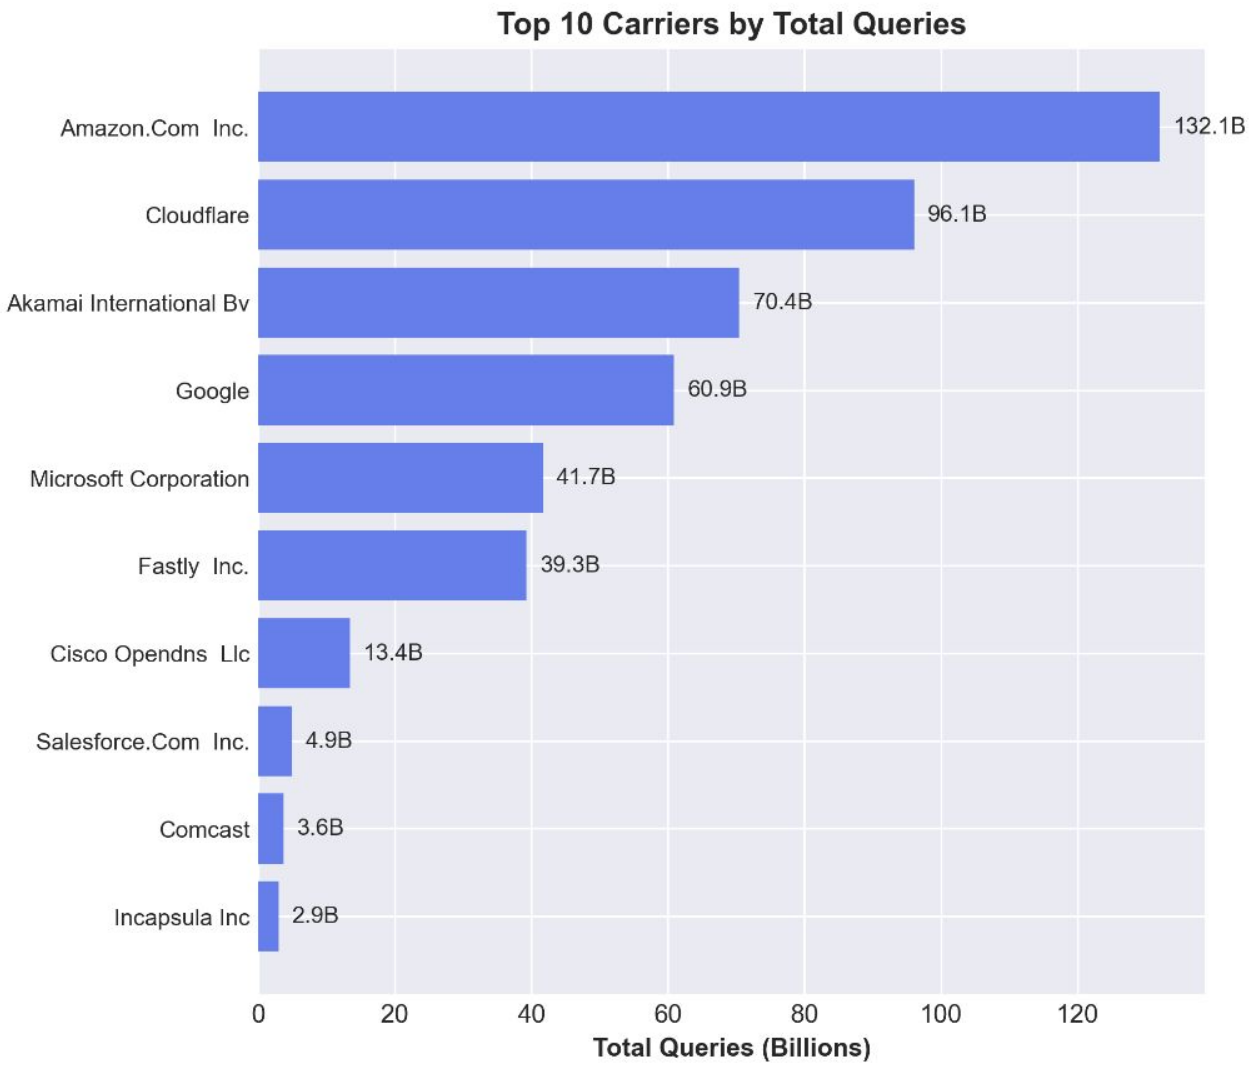

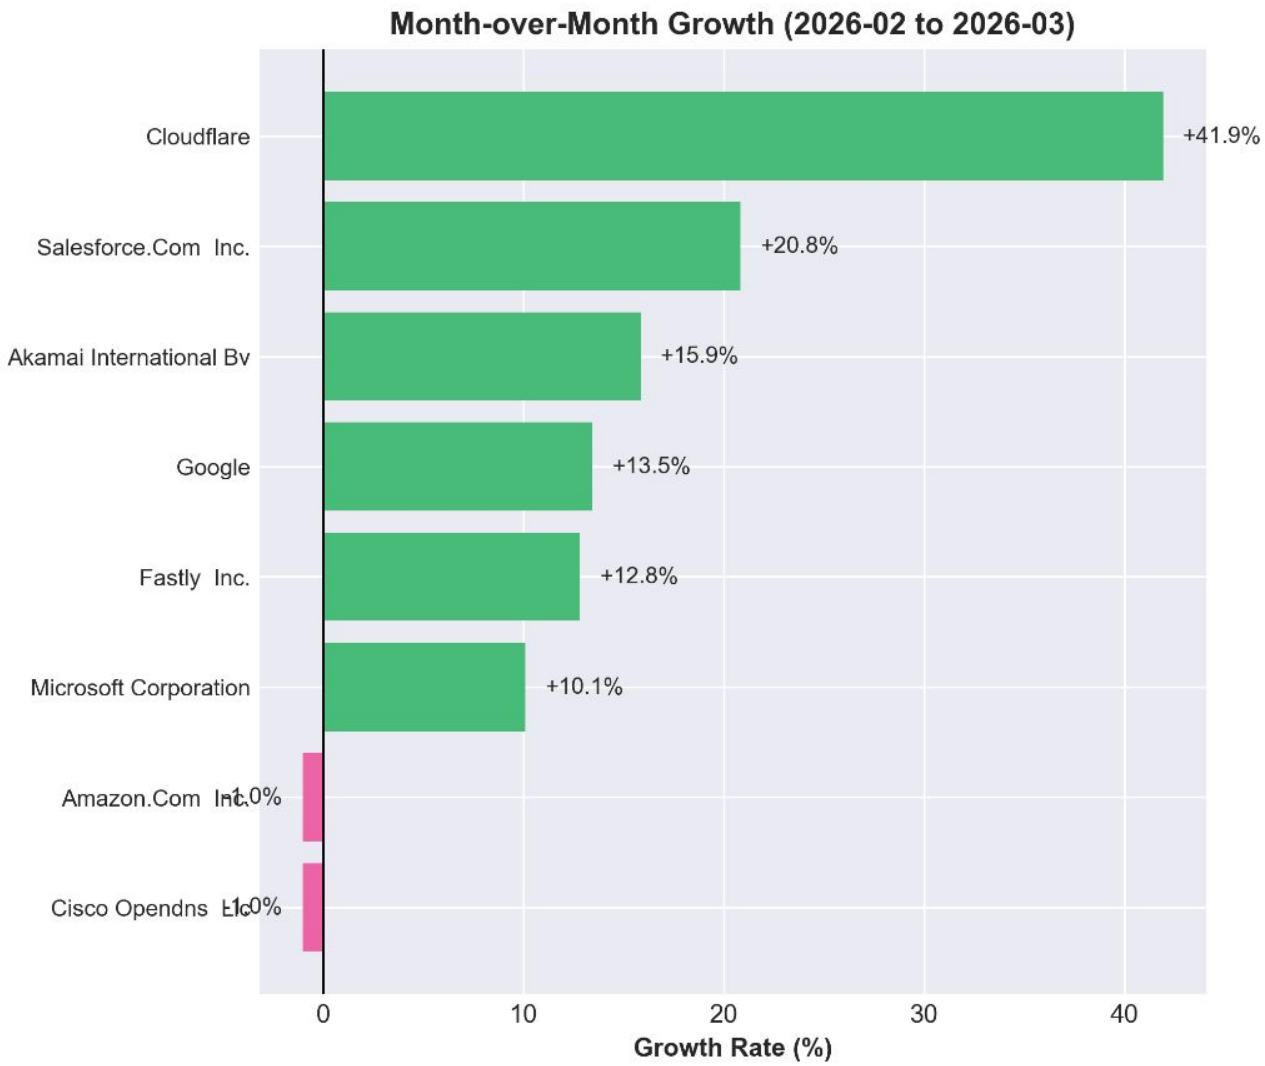

Actionable next steps

Evaluate Caching Strategy

Evaluate caching on corporate resolver to **reduce Query loads**.

Optimize GSLB Design

Review results for **optimal** Global Load balancers (**GSLB**) **design**.

TRAFFIC TRENDS & CHARACTERIZATION

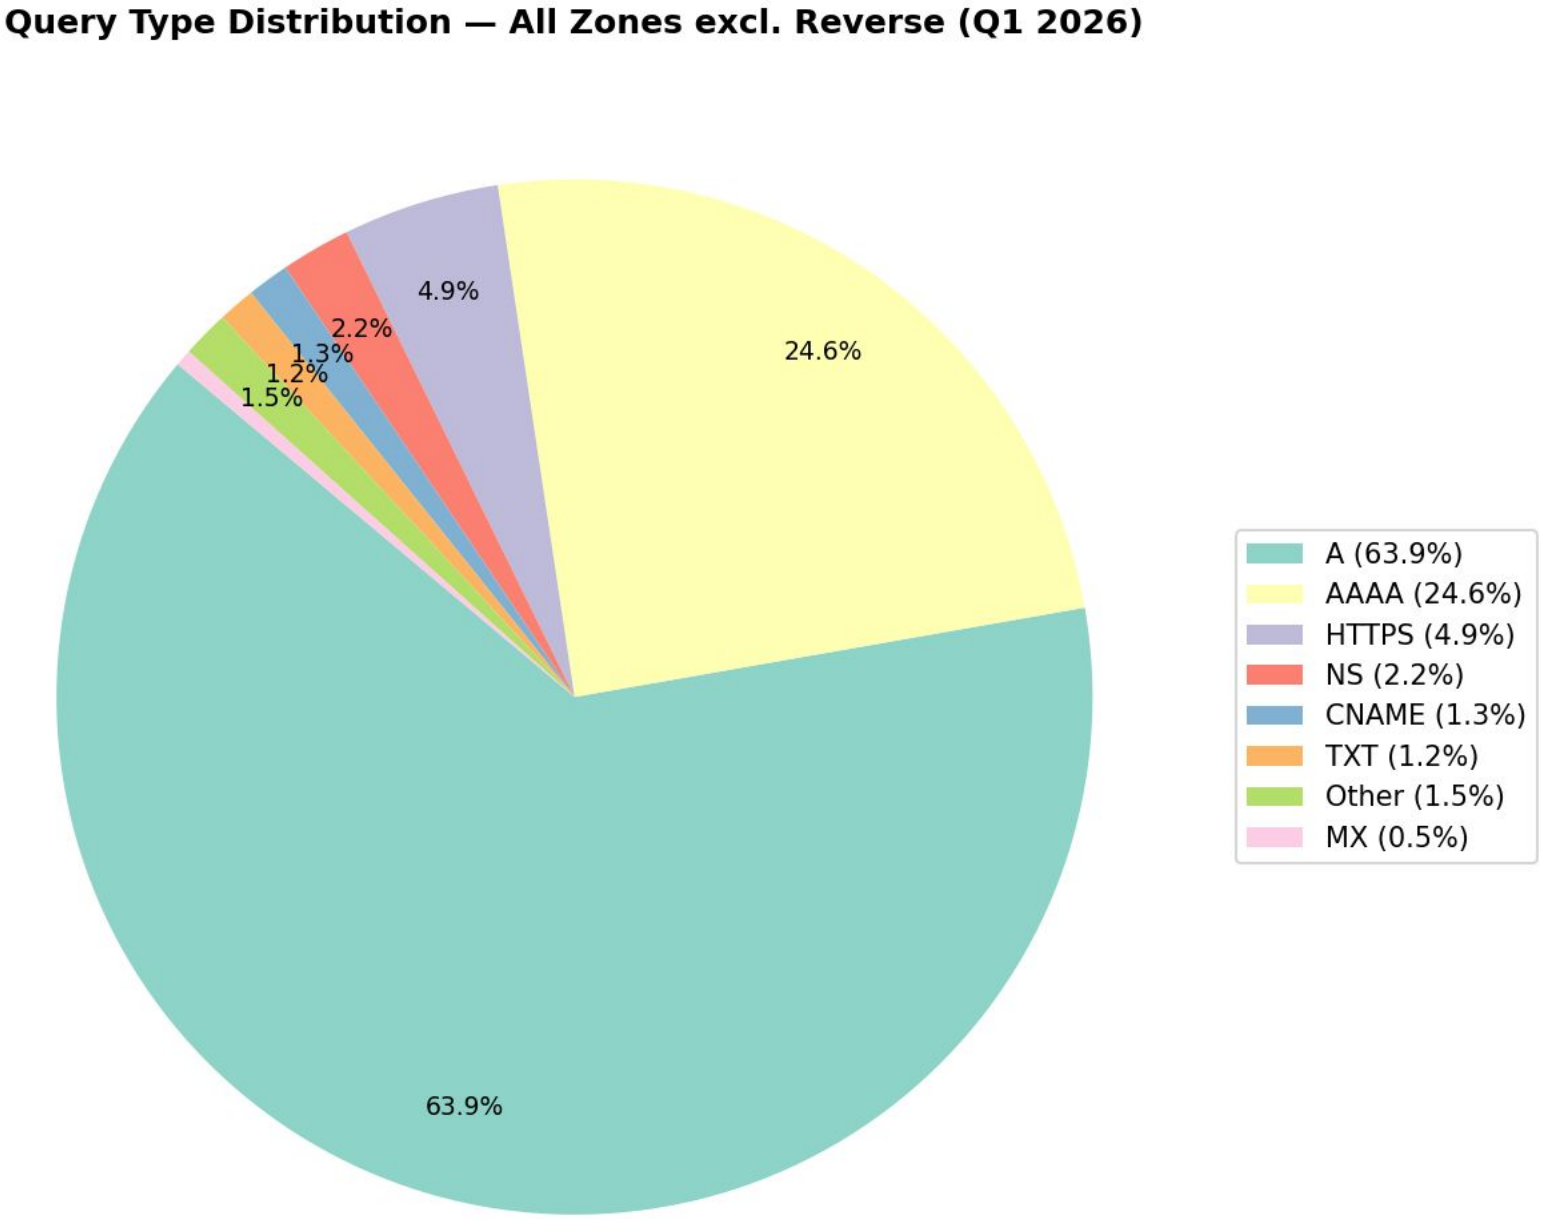
Query Type Distribution — All Zones excl. Reverse (Q1 2026)
4.9%
2.2%
1.3%
1.2%
1.5%
24.6%
63.9%
A (63.9%)
AAAA (24.6%)
HTTPS (4.9%)
NS (2.2%)
CNAME (1.3%)
TXT (1.2%)
Other (1.5%)
MX (0.5%)

Record Type	Response Count	Pct
A	620,153,886,005	63.463%
AAAA	238,254,747,116	24.382%
HTTPS	47,252,122,921	4.835%
NS	21,275,173,054	2.177%
CNAME	12,709,721,465	1.301%
TXT	11,228,906,802	1.149%
OTHER	9,686,359,650	0.991%
PTR	7,305,803,932	0.748%
MX	5,019,777,403	0.514%
DNSKEY	3,074,036,924	0.315%
SOA	810,694,030	0.083%
SRV	130,062,070	0.013%
RRSIG	117,499,326	0.012%
SPF	78,079,082	0.008%
ANY	62,580,219	0.006%
HINFO	18,582,380	0.002%
NSEC3	9,003,183	0.001%

Record Type	Response Count	Pct
NSEC	1,859,756	0.0%
A6	1,069,152	0.0%
SSHFP	756,919	0.0%
NAPTR	660,548	0.0%
SVCB	477,890	0.0%
TSIG	177,248	0.0%
CERT	139,799	0.0%
LOC	95,689	0.0%
NSEC3PARAM	51,267	0.0%
RP	30,992	0.0%
IPSECKEY	26,001	0.0%
DLV	24,042	0.0%
TA	23,787	0.0%
TKEY	19,678	0.0%
MF	10,453	0.0%

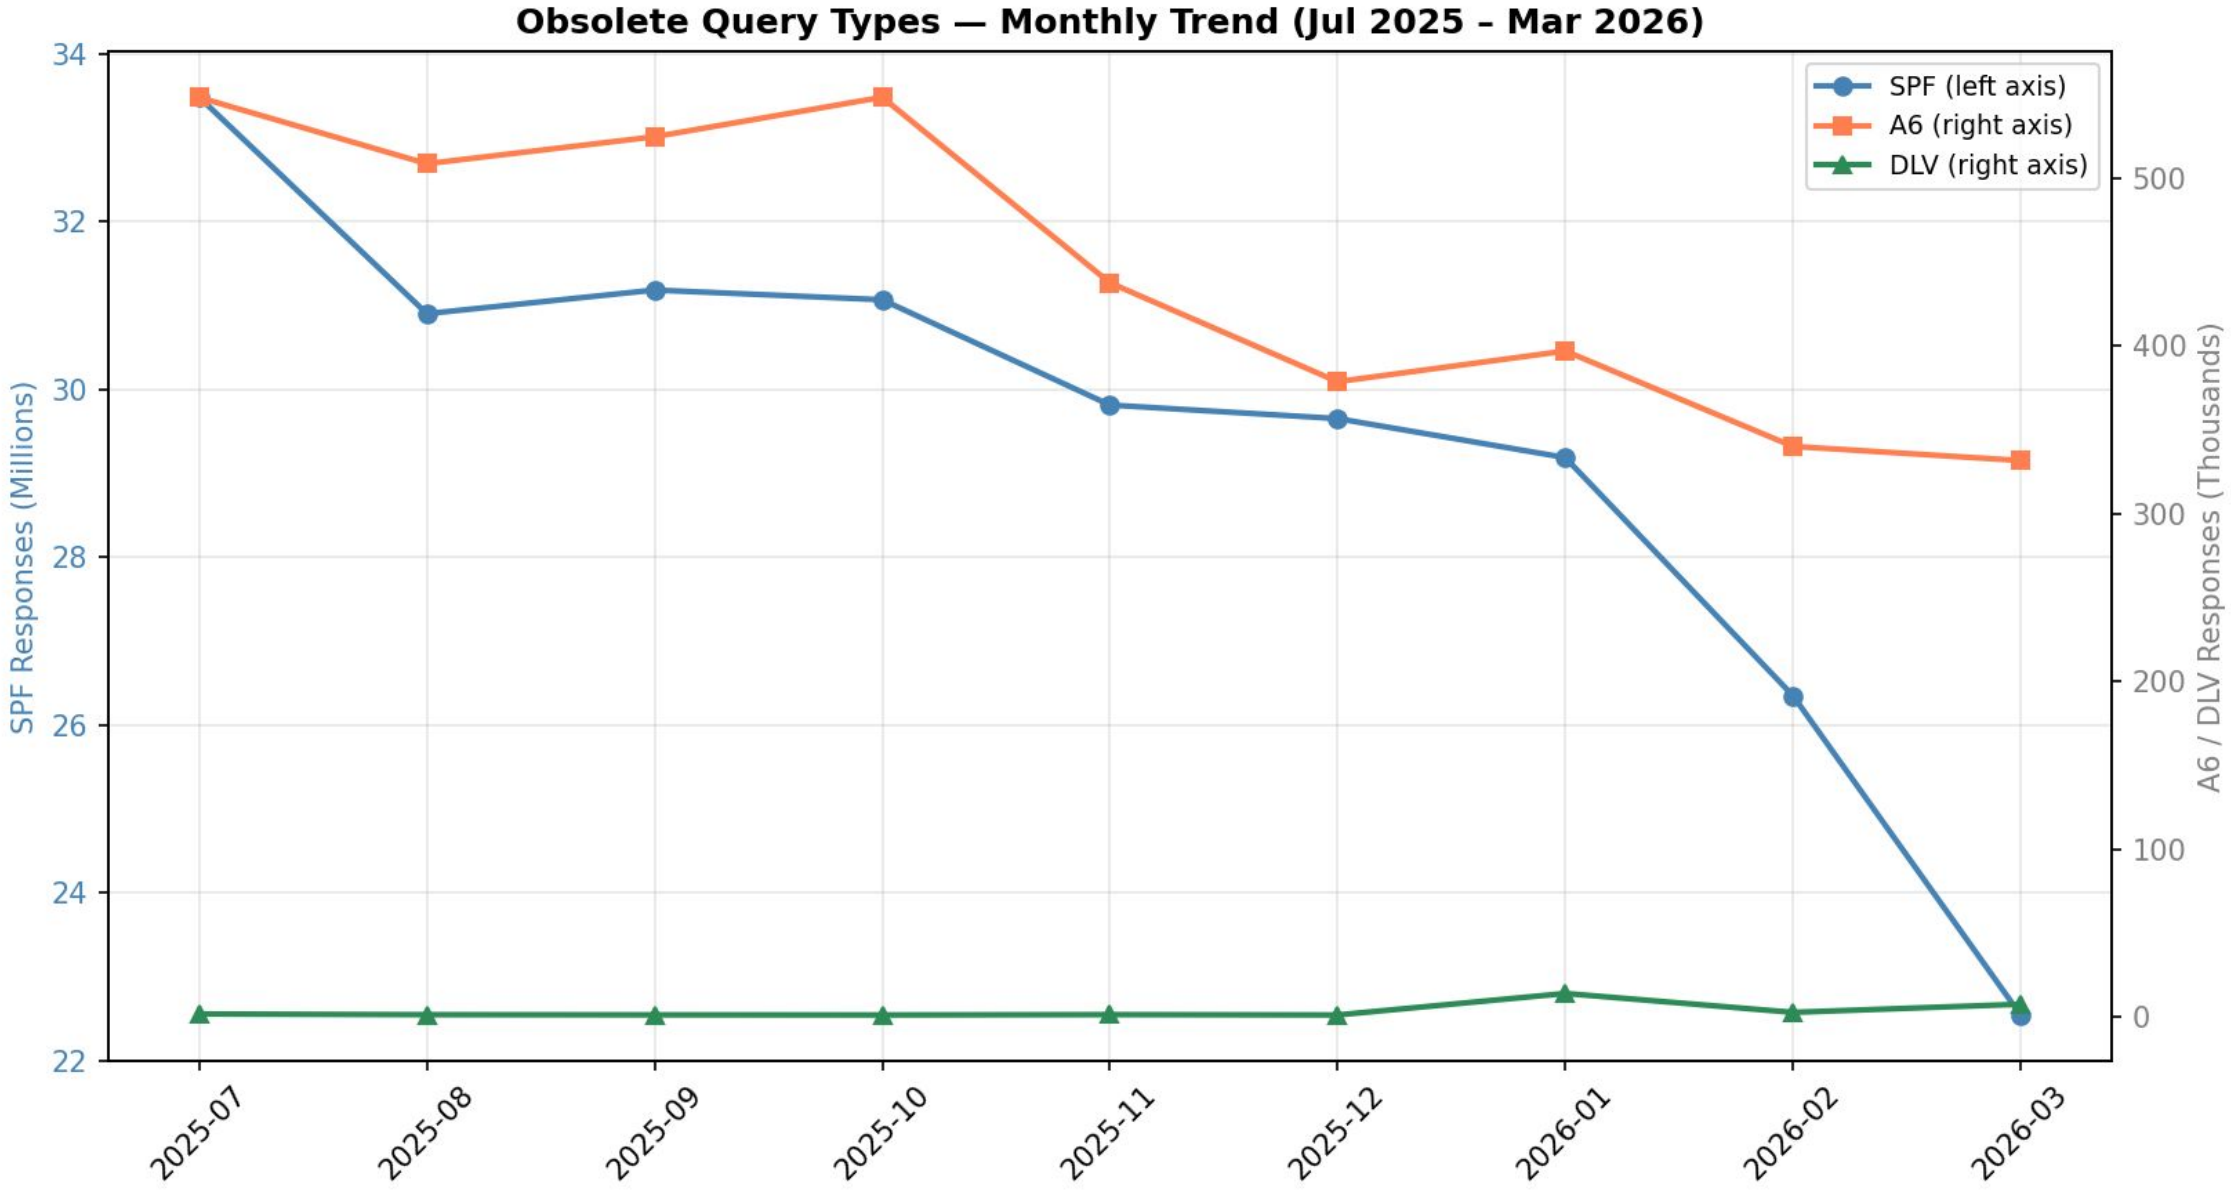
Obsolete Query Types — Monthly Trend (Jul 2025 – Mar 2026)
SPF (left axis)
A6 (right axis)
DLV (right axis)
SPF Responses (Millions)
A6 / DLV Responses (Thousands)
34
32
30
28
26
24
22
500
400
300
200
100
0
2025-07
2025-08
2025-09
2025-10
2025-11
2025-12
2026-01
2026-02
2026-03

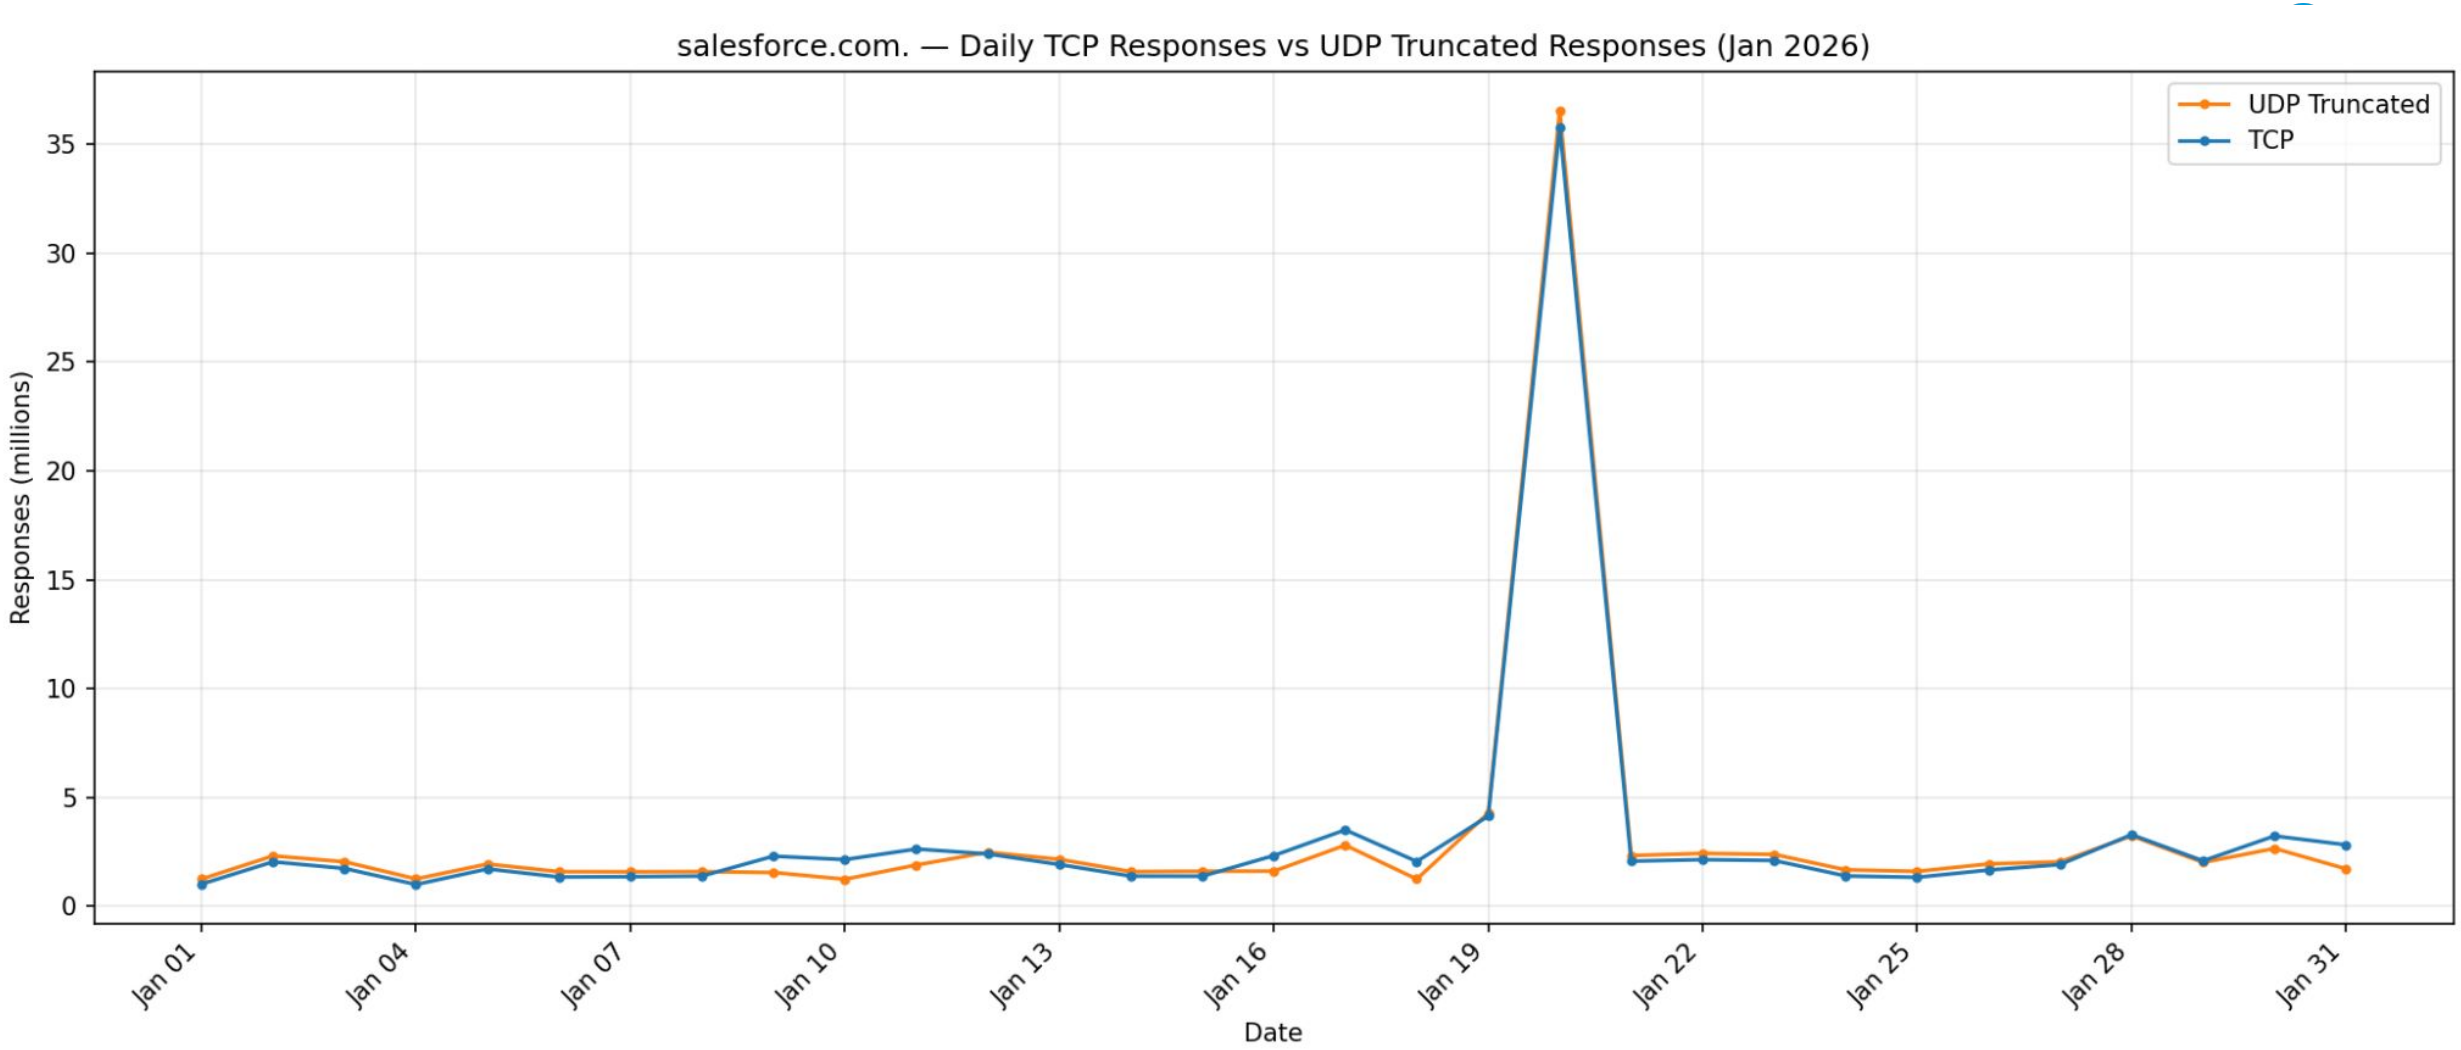
salesforce.com. — Daily TCP Responses vs UDP Truncated Responses (Jan 2026)
UDP Truncated
TCP
Responses (millions)
0
5
10
15
20
25
30
35
Jan 01
Jan 04
Jan 07
Jan 10
Jan 13
Jan 16
Jan 19
Jan 22
Jan 25
Jan 28
Jan 31
Date

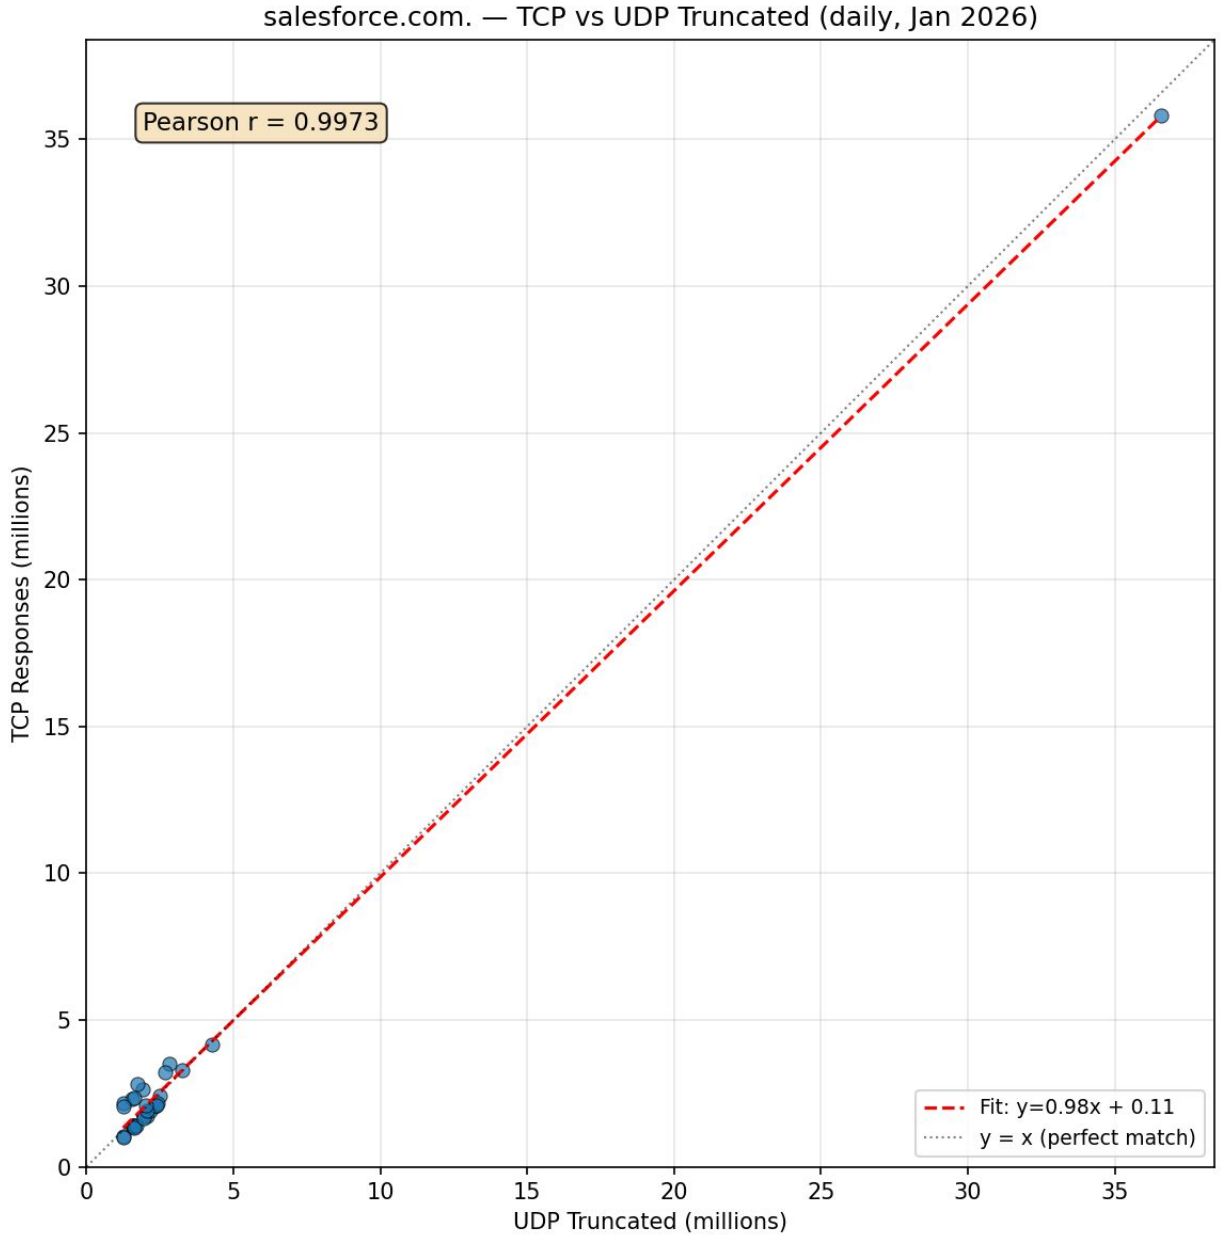
salesforce.com. — TCP vs UDP Truncated (daily, Jan 2026)
Pearson r = 0.9973
TCP Responses (millions)
UDP Truncated (millions)
0
5
10
15
20
25
30
35
Fit: y=0.98x + 0.11
y = x (perfect match)

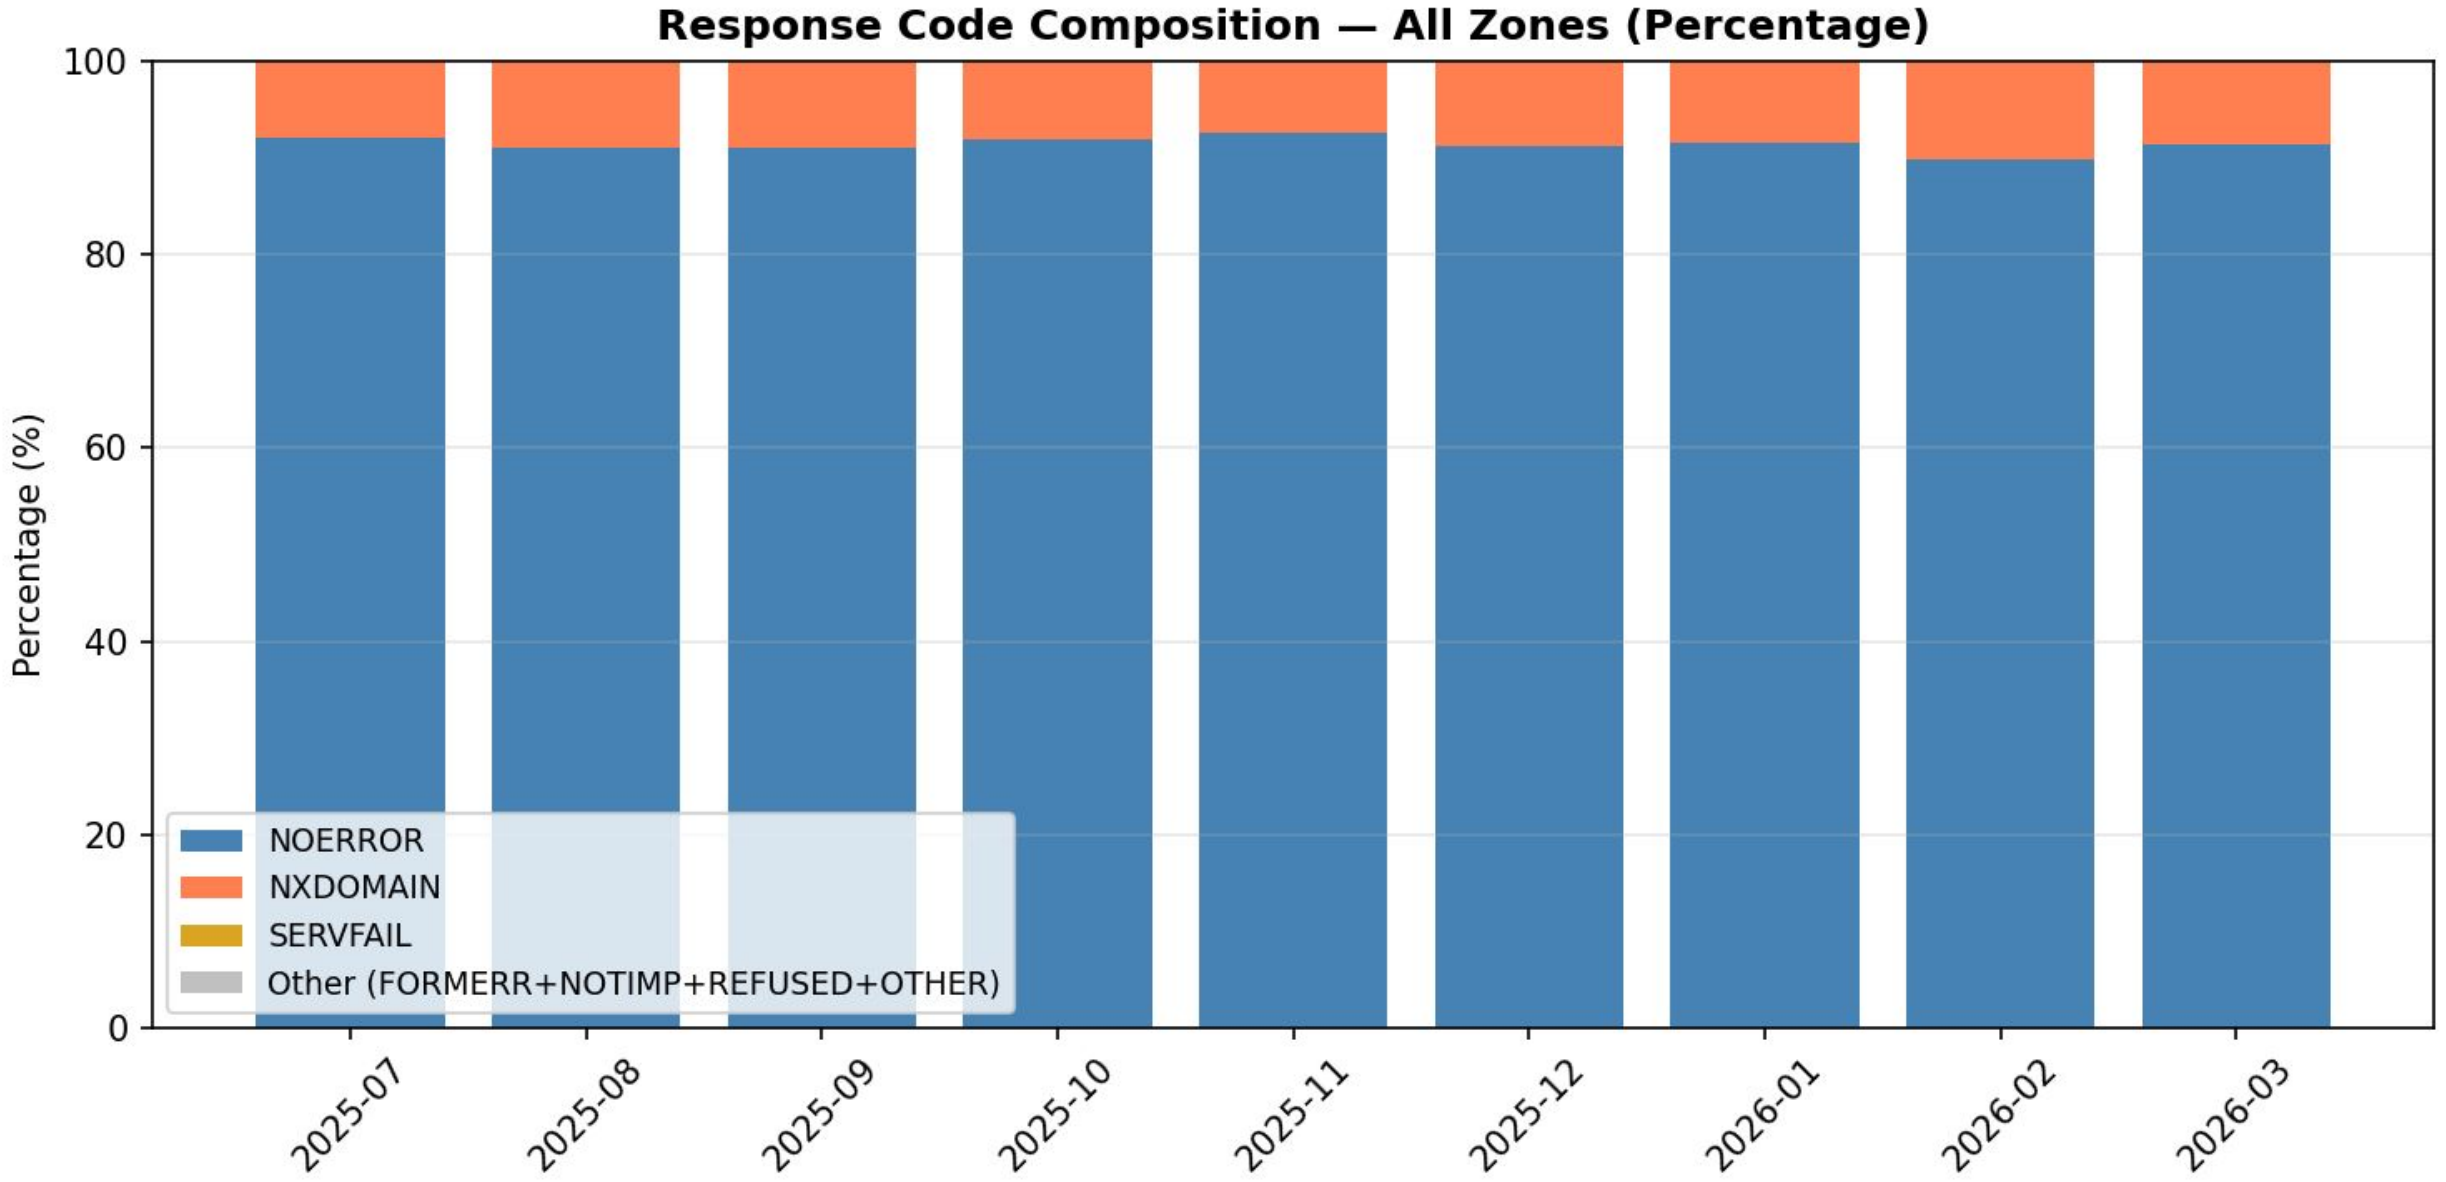
Response Code Composition — All Zones (Percentage)
Percentage (%)
100
80
60
40
20
0
NOERROR
NXDOMAIN
SERVFAIL
Other (FORMERR+NOTIMP+REFUSED+OTHER)
2025-07
2025-08
2025-09
2025-10
2025-11
2025-12
2026-01
2026-02
2026-03

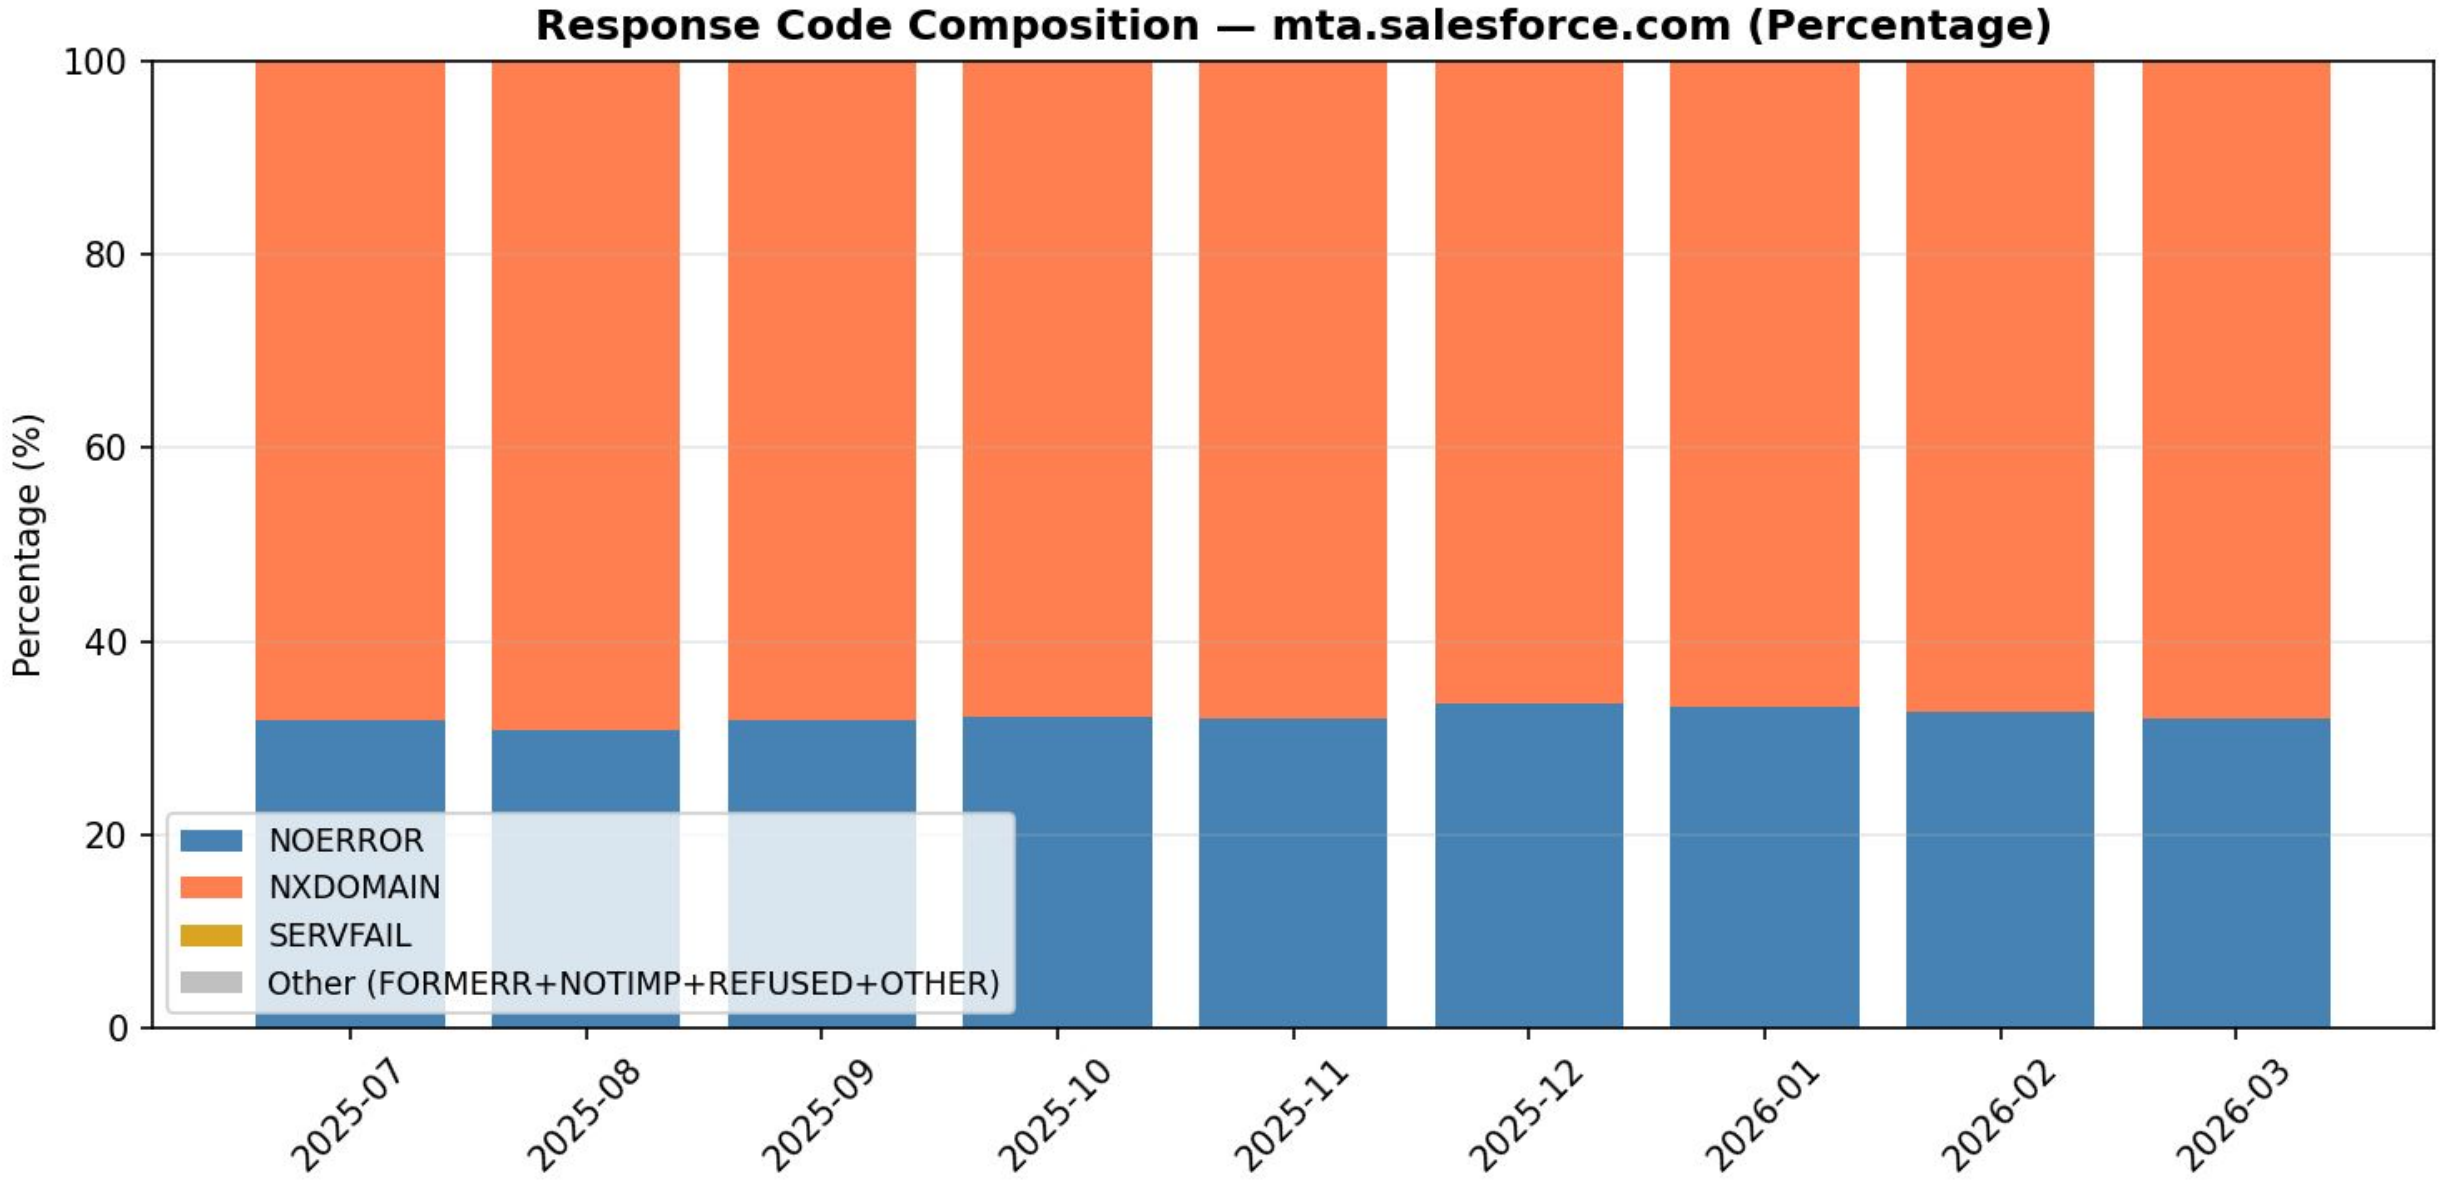
Response Code Composition — mta.salesforce.com (Percentage)
Percentage (%)
100
80
60
40
20
0
NOERROR
NXDOMAIN
SERVFAIL
Other (FORMERR+NOTIMP+REFUSED+OTHER)
2025-07
2025-08
2025-09
2025-10
2025-11
2025-12
2026-01
2026-02
2026-03

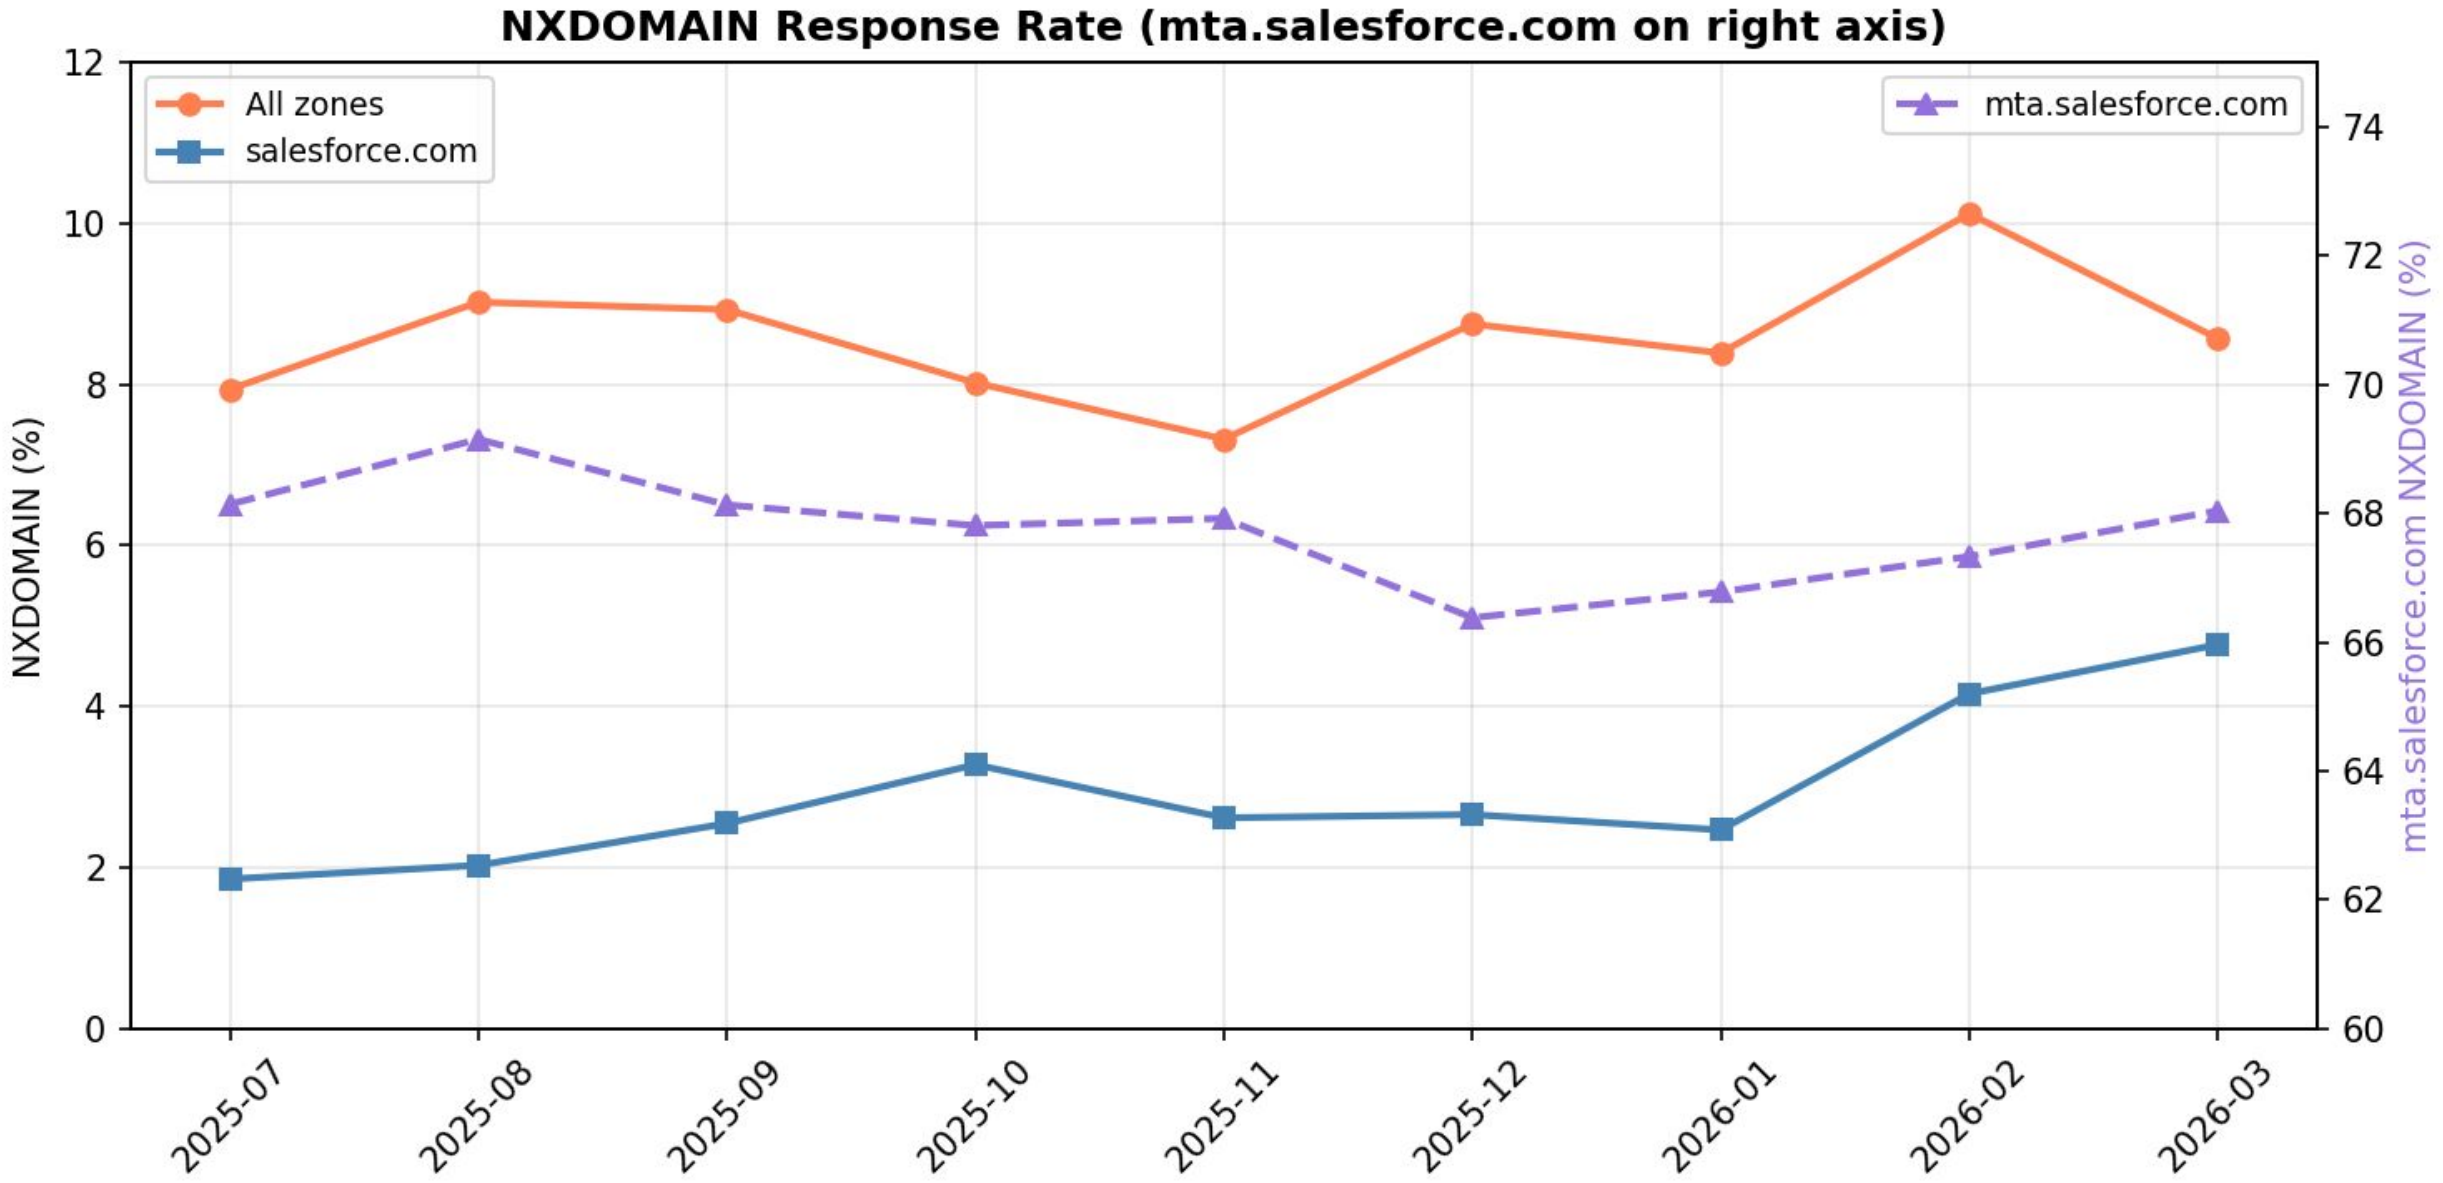
NXDOMAIN Response Rate (mta.salesforce.com on right axis)
All zones
salesforce.com
mta.salesforce.com
NXDOMAIN (%)
mta.salesforce.com NXDOMAIN (%)
0
2
4
6
8
10
12
60
62
64
66
68
70
72
74
2025-07
2025-08
2025-09
2025-10
2025-11
2025-12
2026-01
2026-02
2026-03

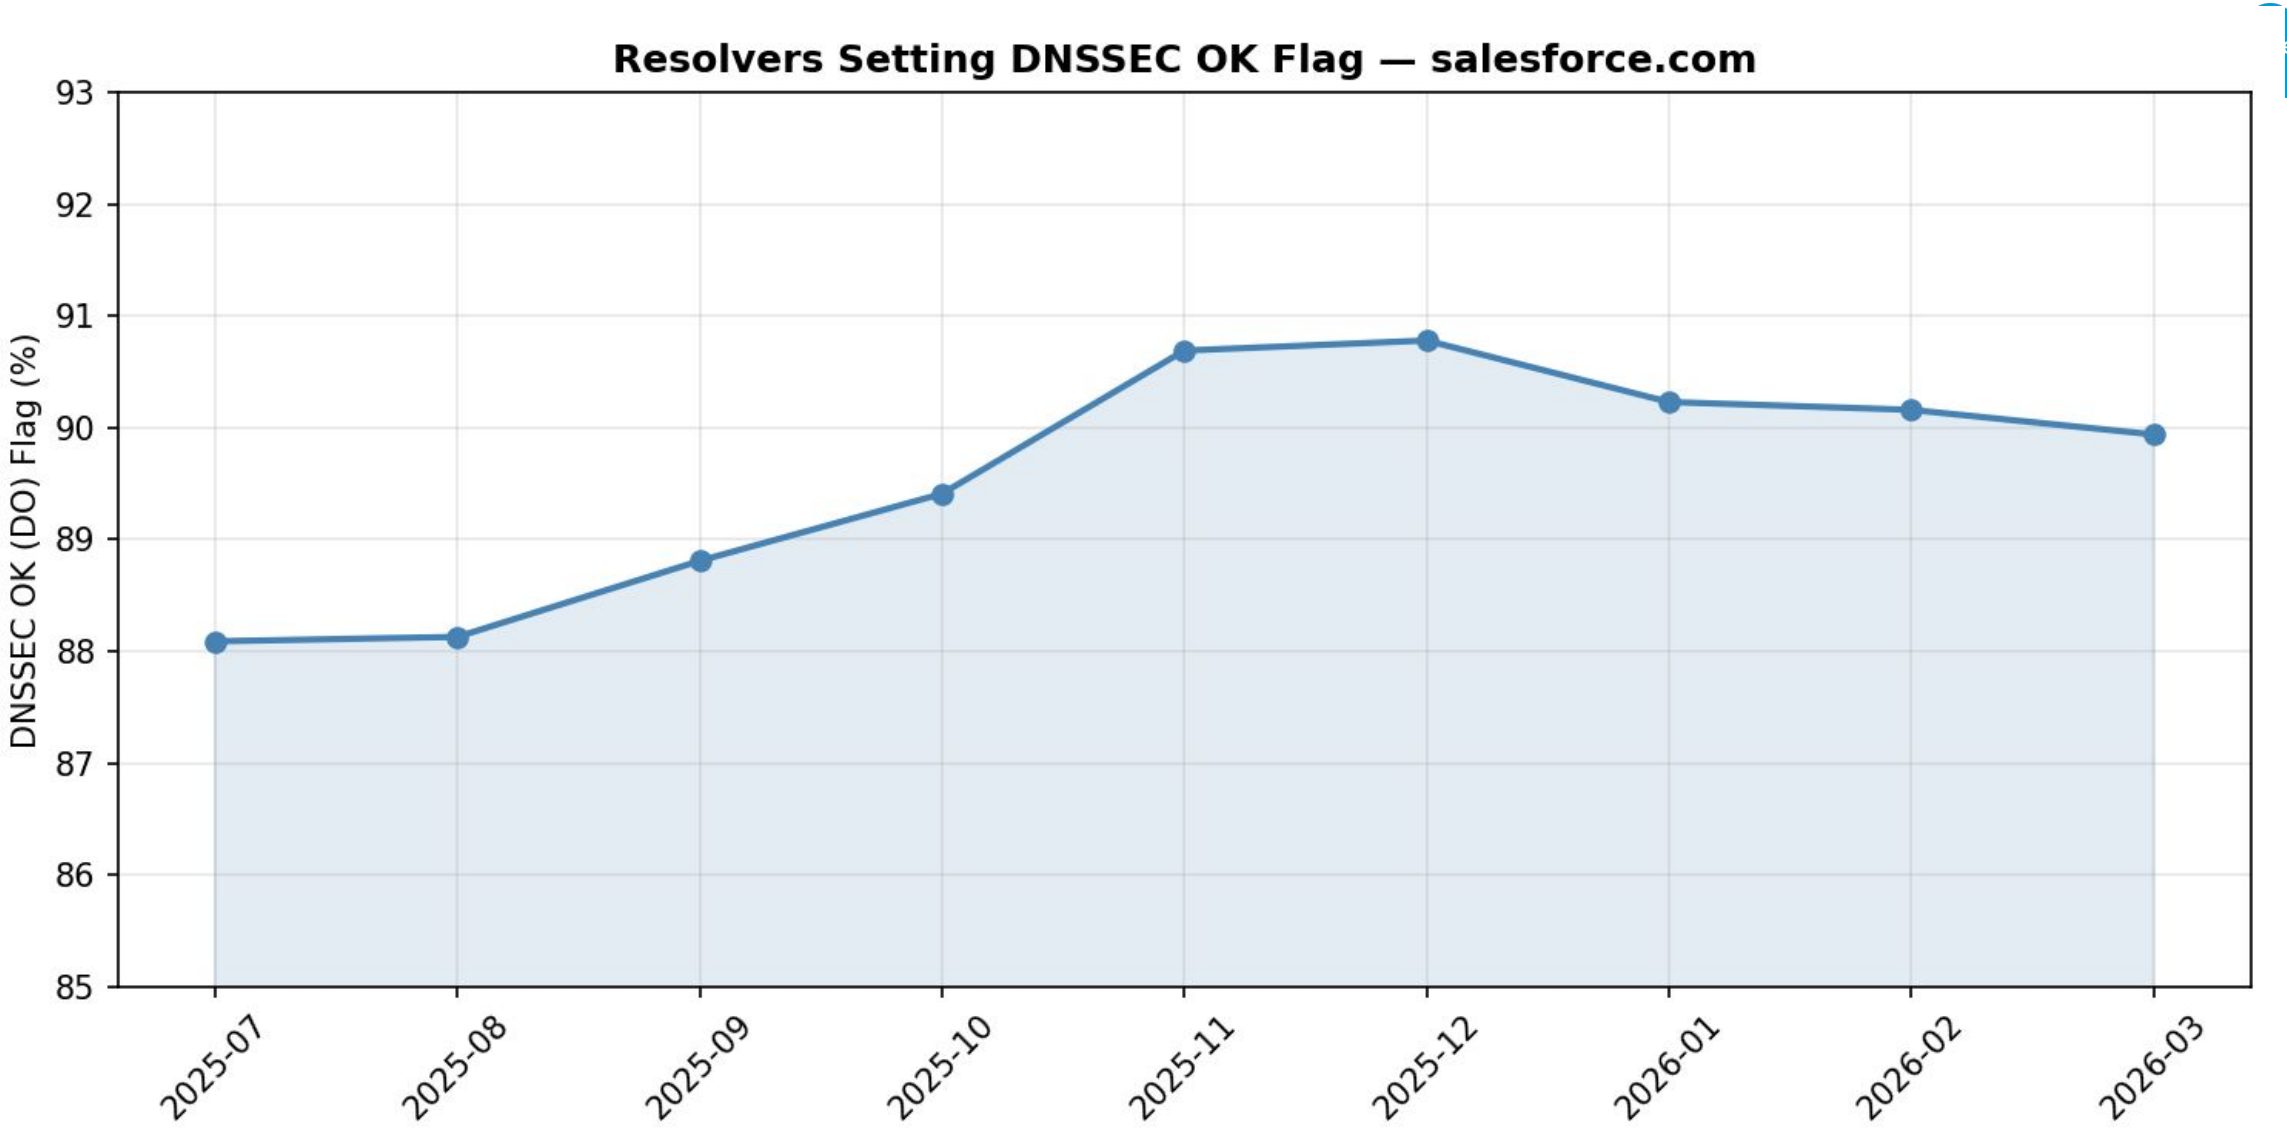
Resolvers Setting DNSSEC OK Flag — salesforce.com
DNSSEC OK (DO) Flag (%)
93
92
91
90
89
88
87
86
85
2025-07
2025-08
2025-09
2025-10
2025-11
2025-12
2026-01
2026-02
2026-03

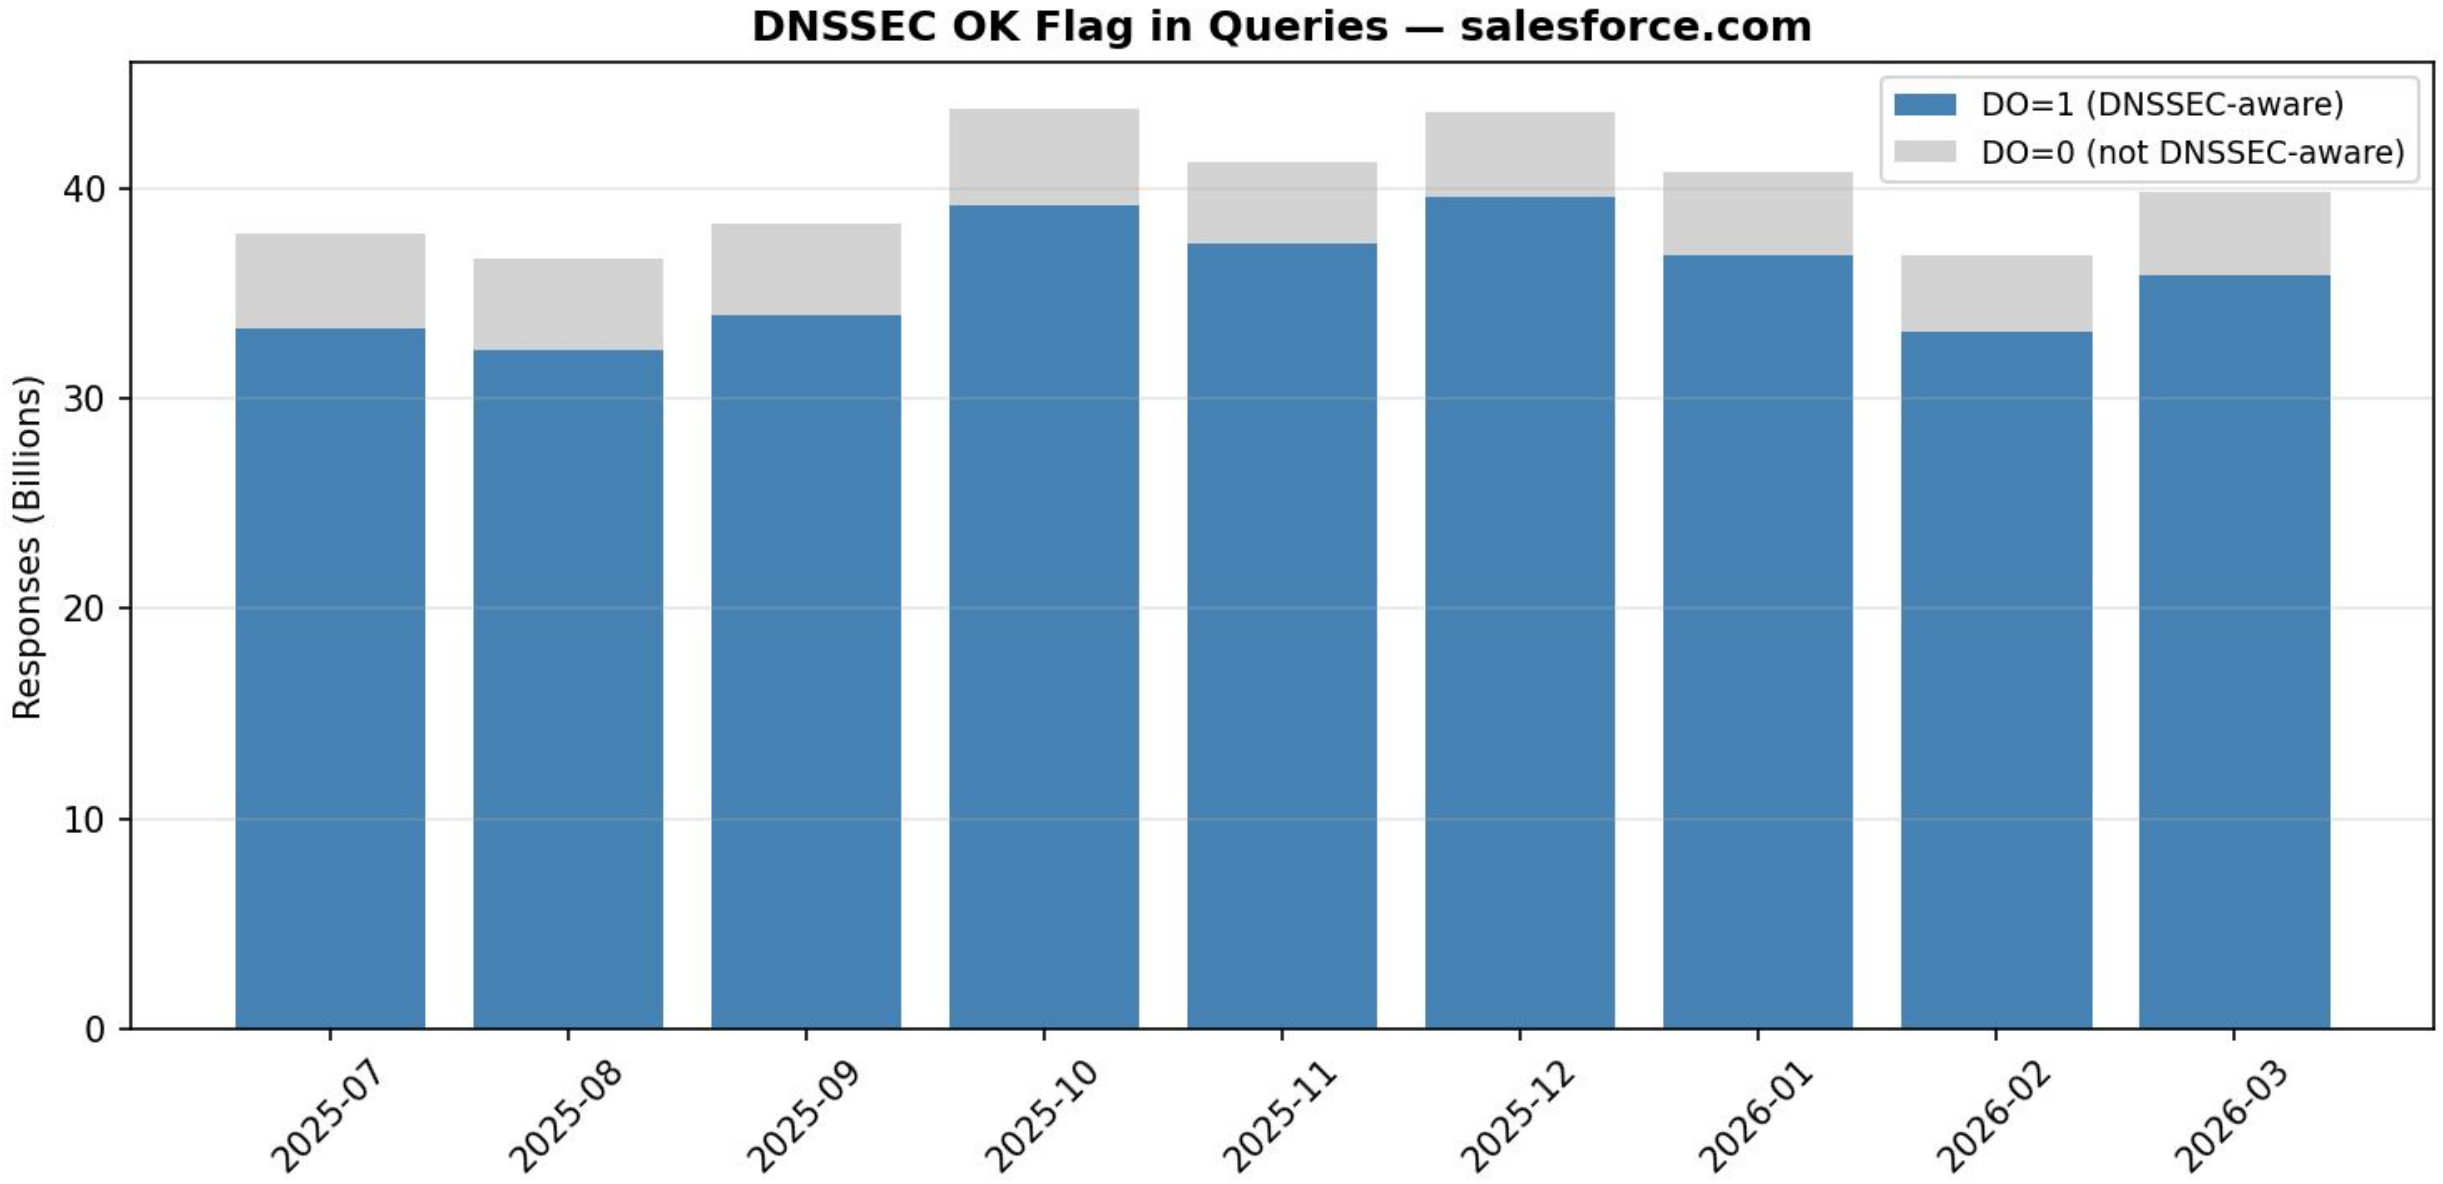
DNSSEC OK Flag in Queries — salesforce.com
DO=1 (DNSSEC-aware)
DO=0 (not DNSSEC-aware)
Responses (Billions)
0
10
20
30
40
2025-07
2025-08
2025-09
2025-10
2025-11
2025-12
2026-01
2026-02
2026-03

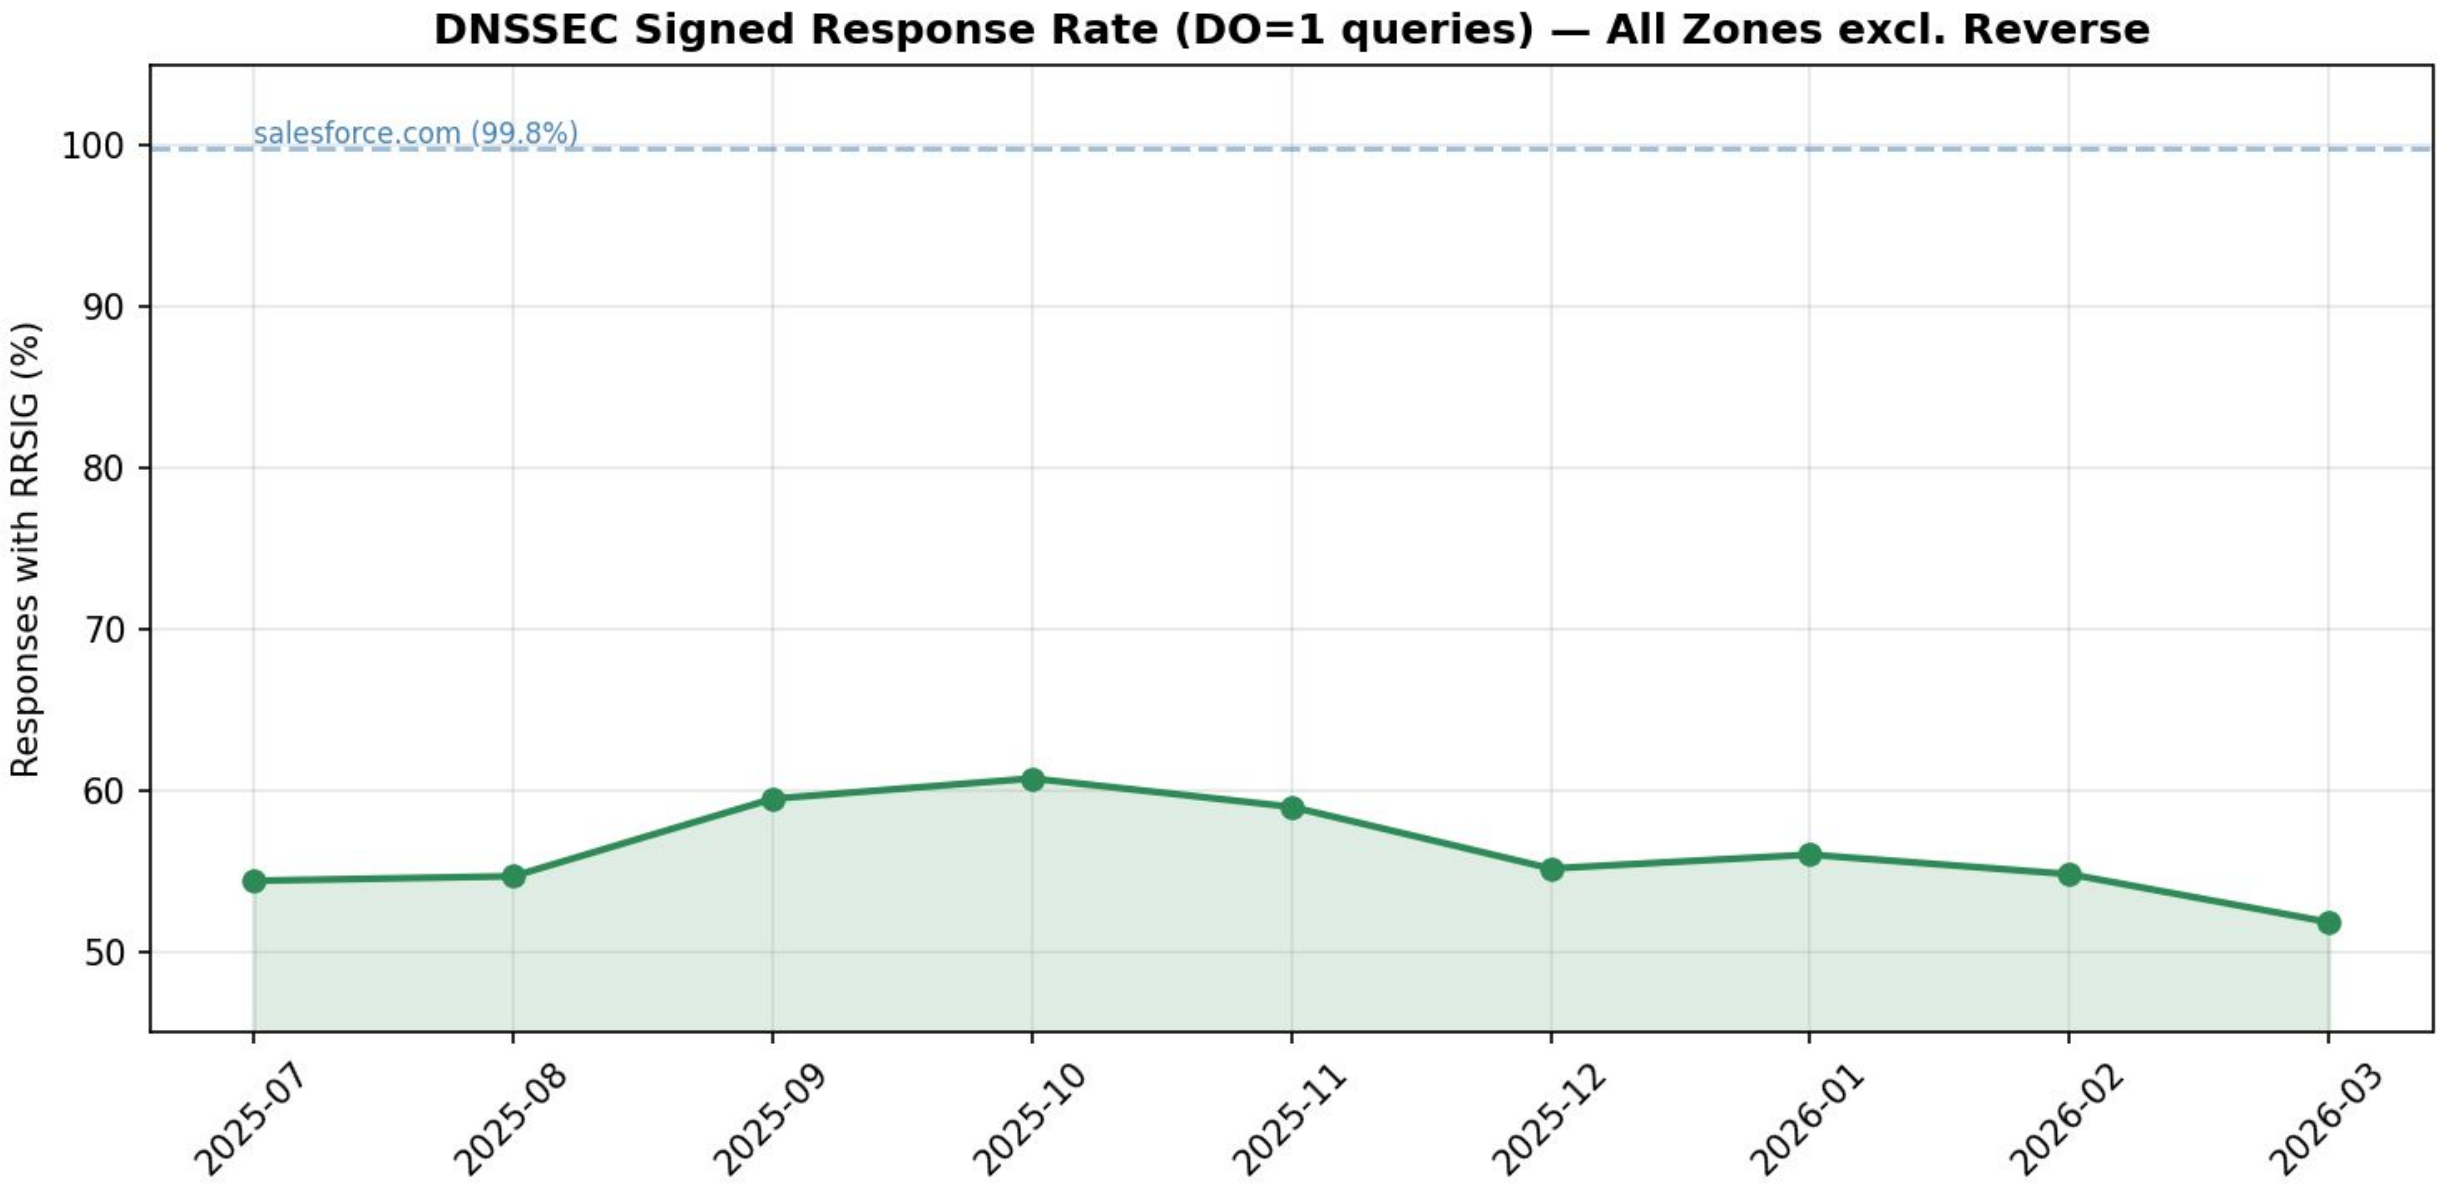
DNSSEC Signed Response Rate (DO=1 queries) — All Zones excl. Reverse
salesforce.com (99.8%)
Responses with RRSIG (%)
100
90
80
70
60
50
2025-07
2025-08
2025-09
2025-10
2025-11
2025-12
2026-01
2026-02
2026-03

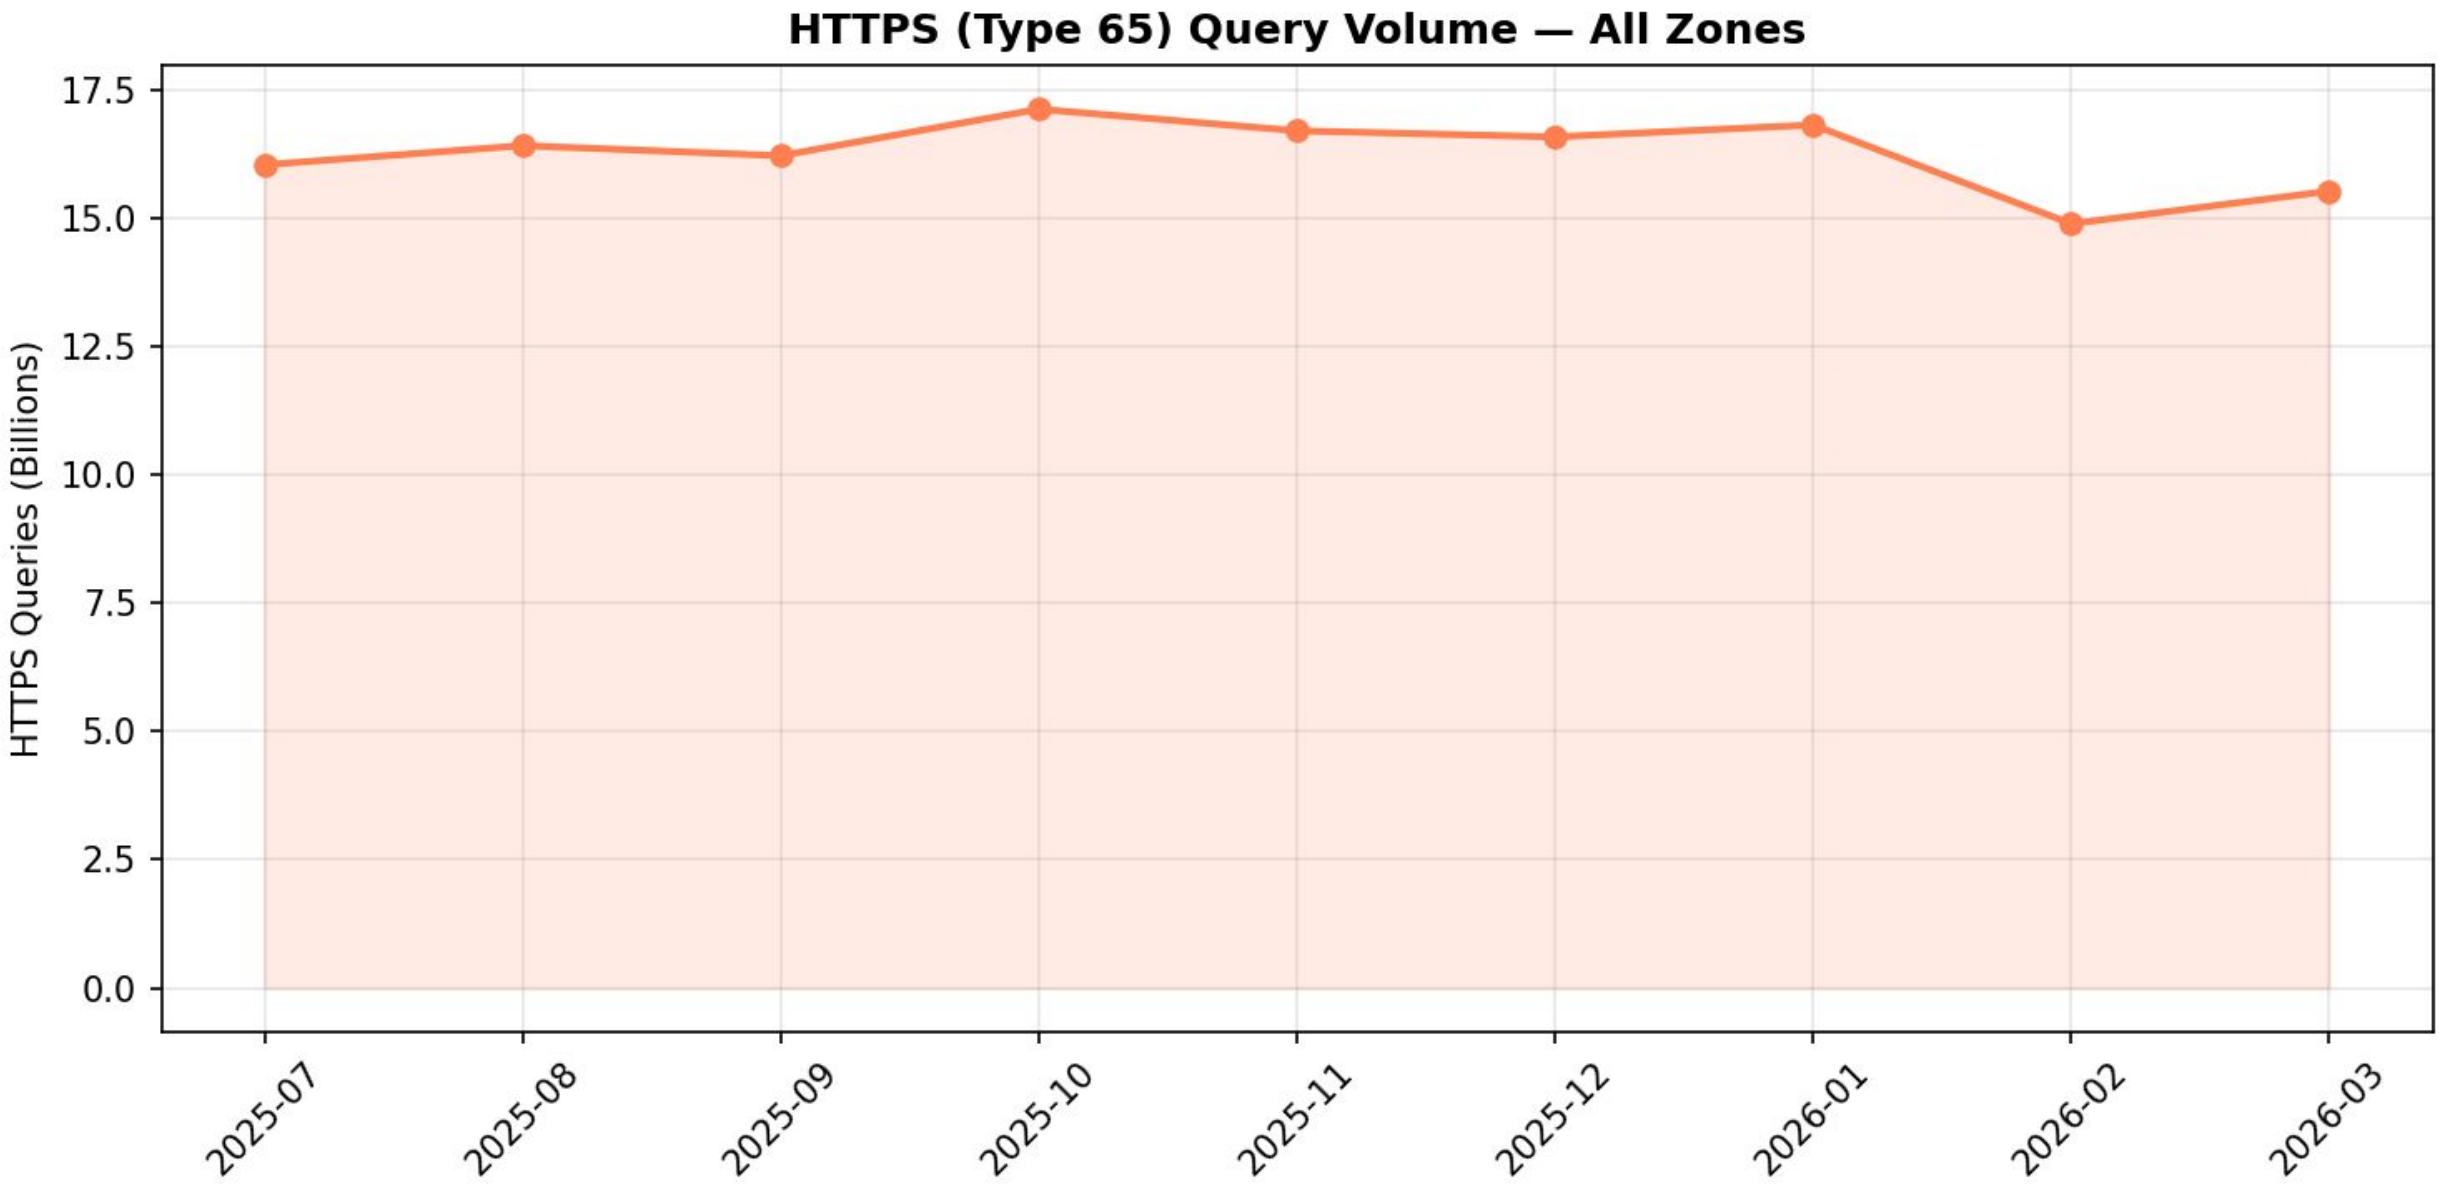
HTTPS (Type 65) Query Volume — All Zones
HTTPS Queries (Billions)
17.5
15.0
12.5
10.0
7.5
5.0
2.5
0.0
2025-07
2025-08
2025-09
2025-10
2025-11
2025-12
2026-01
2026-02
2026-03

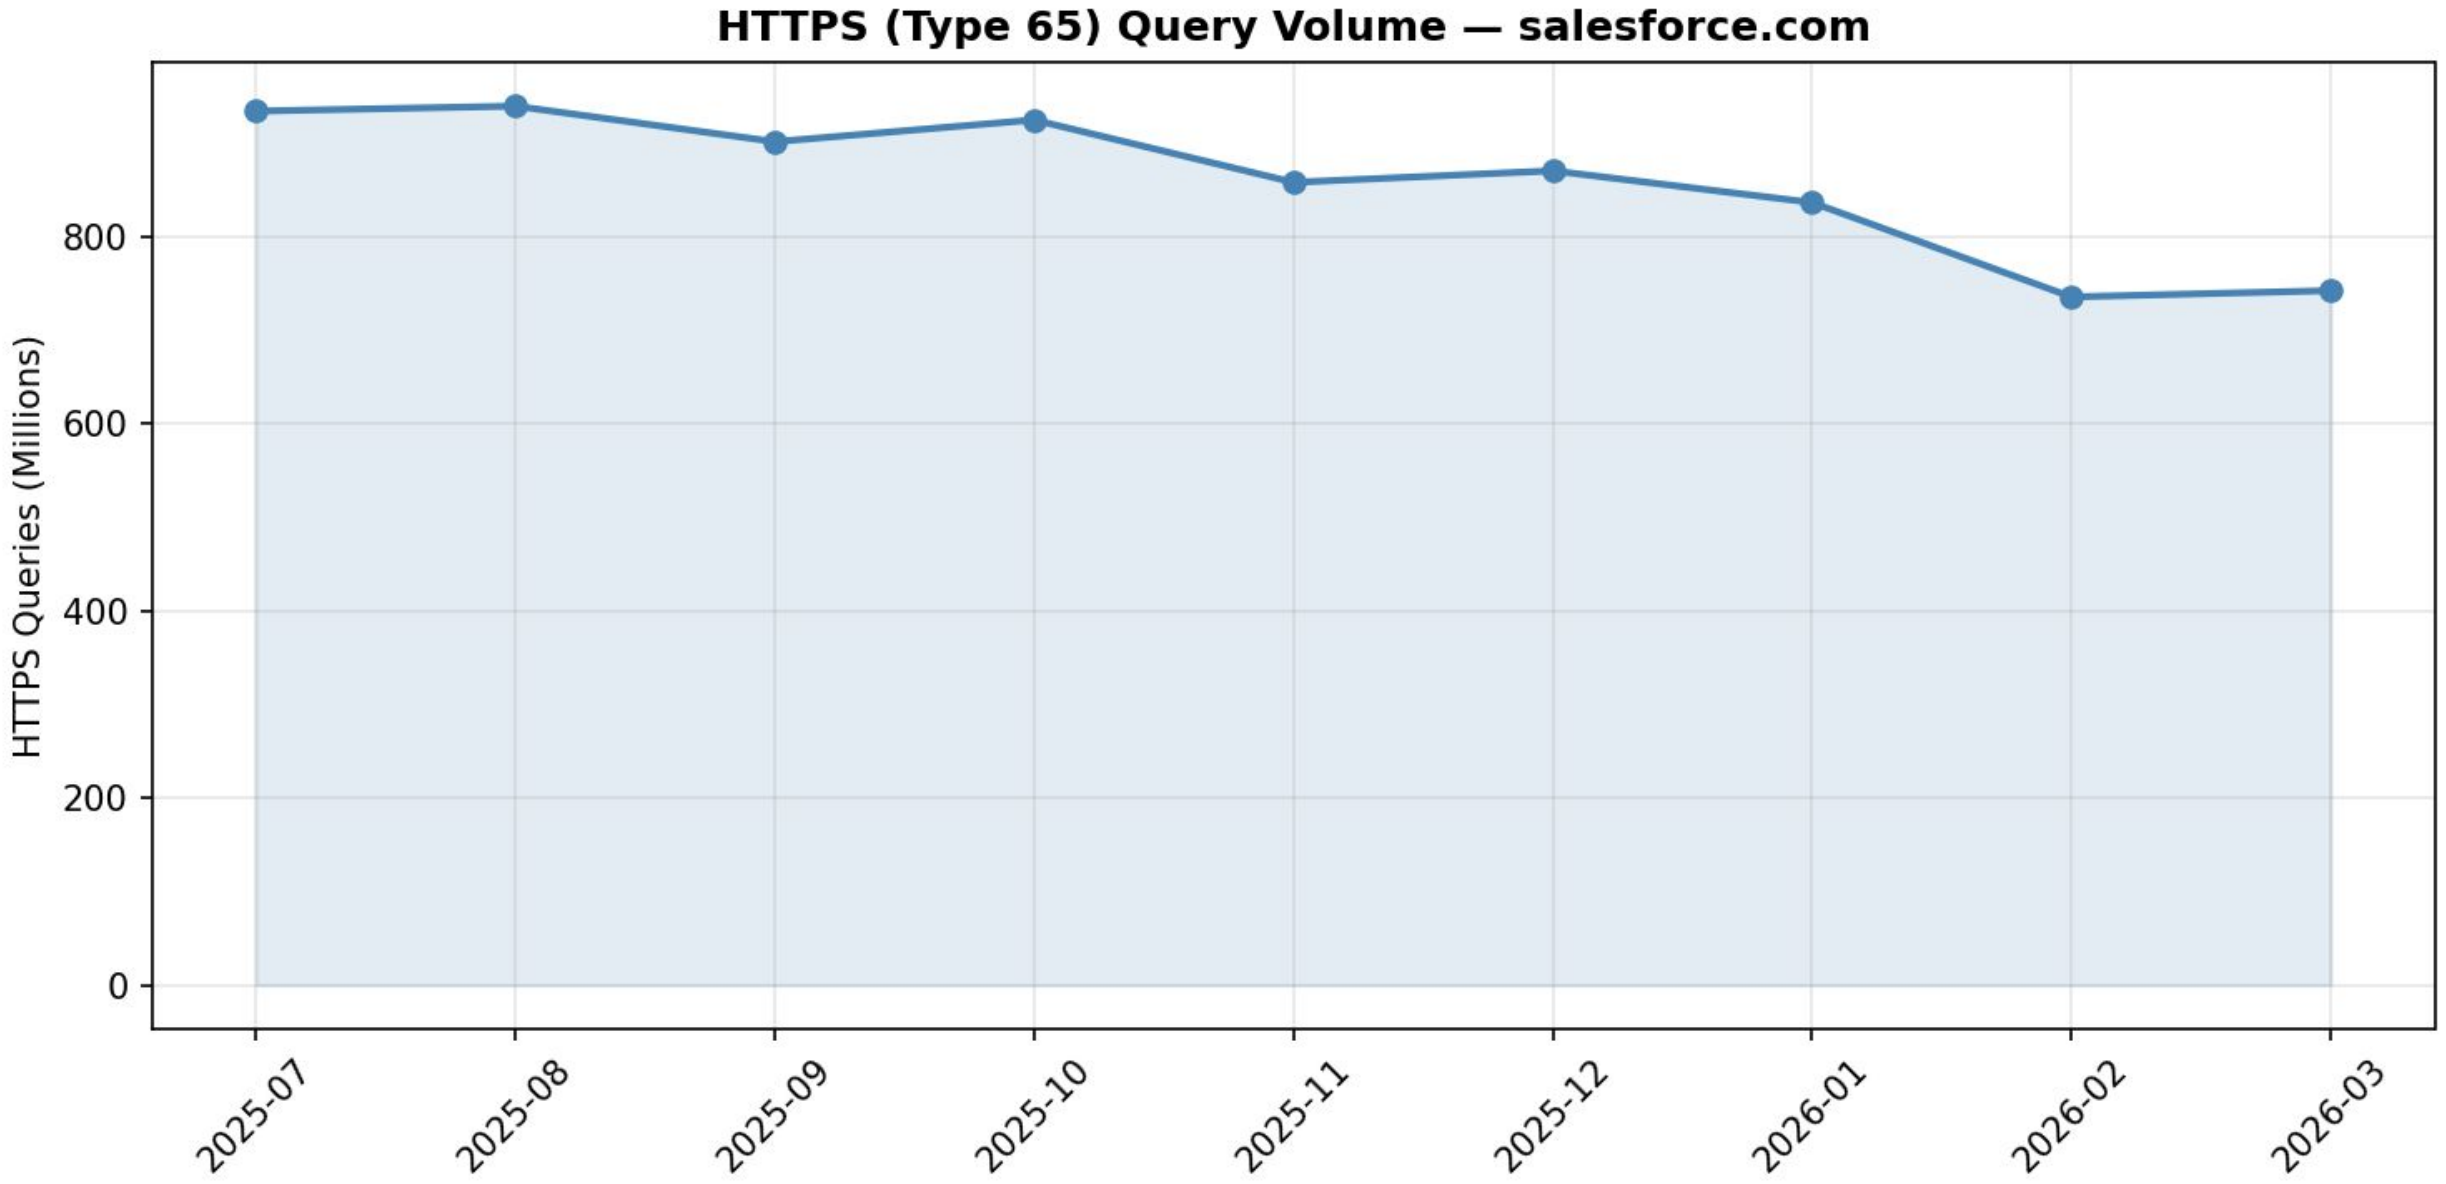
HTTPS (Type 65) Query Volume — salesforce.com
HTTPS Queries (Millions)
0
200
400
600
800
2025-07
2025-08
2025-09
2025-10
2025-11
2025-12
2026-01
2026-02
2026-03

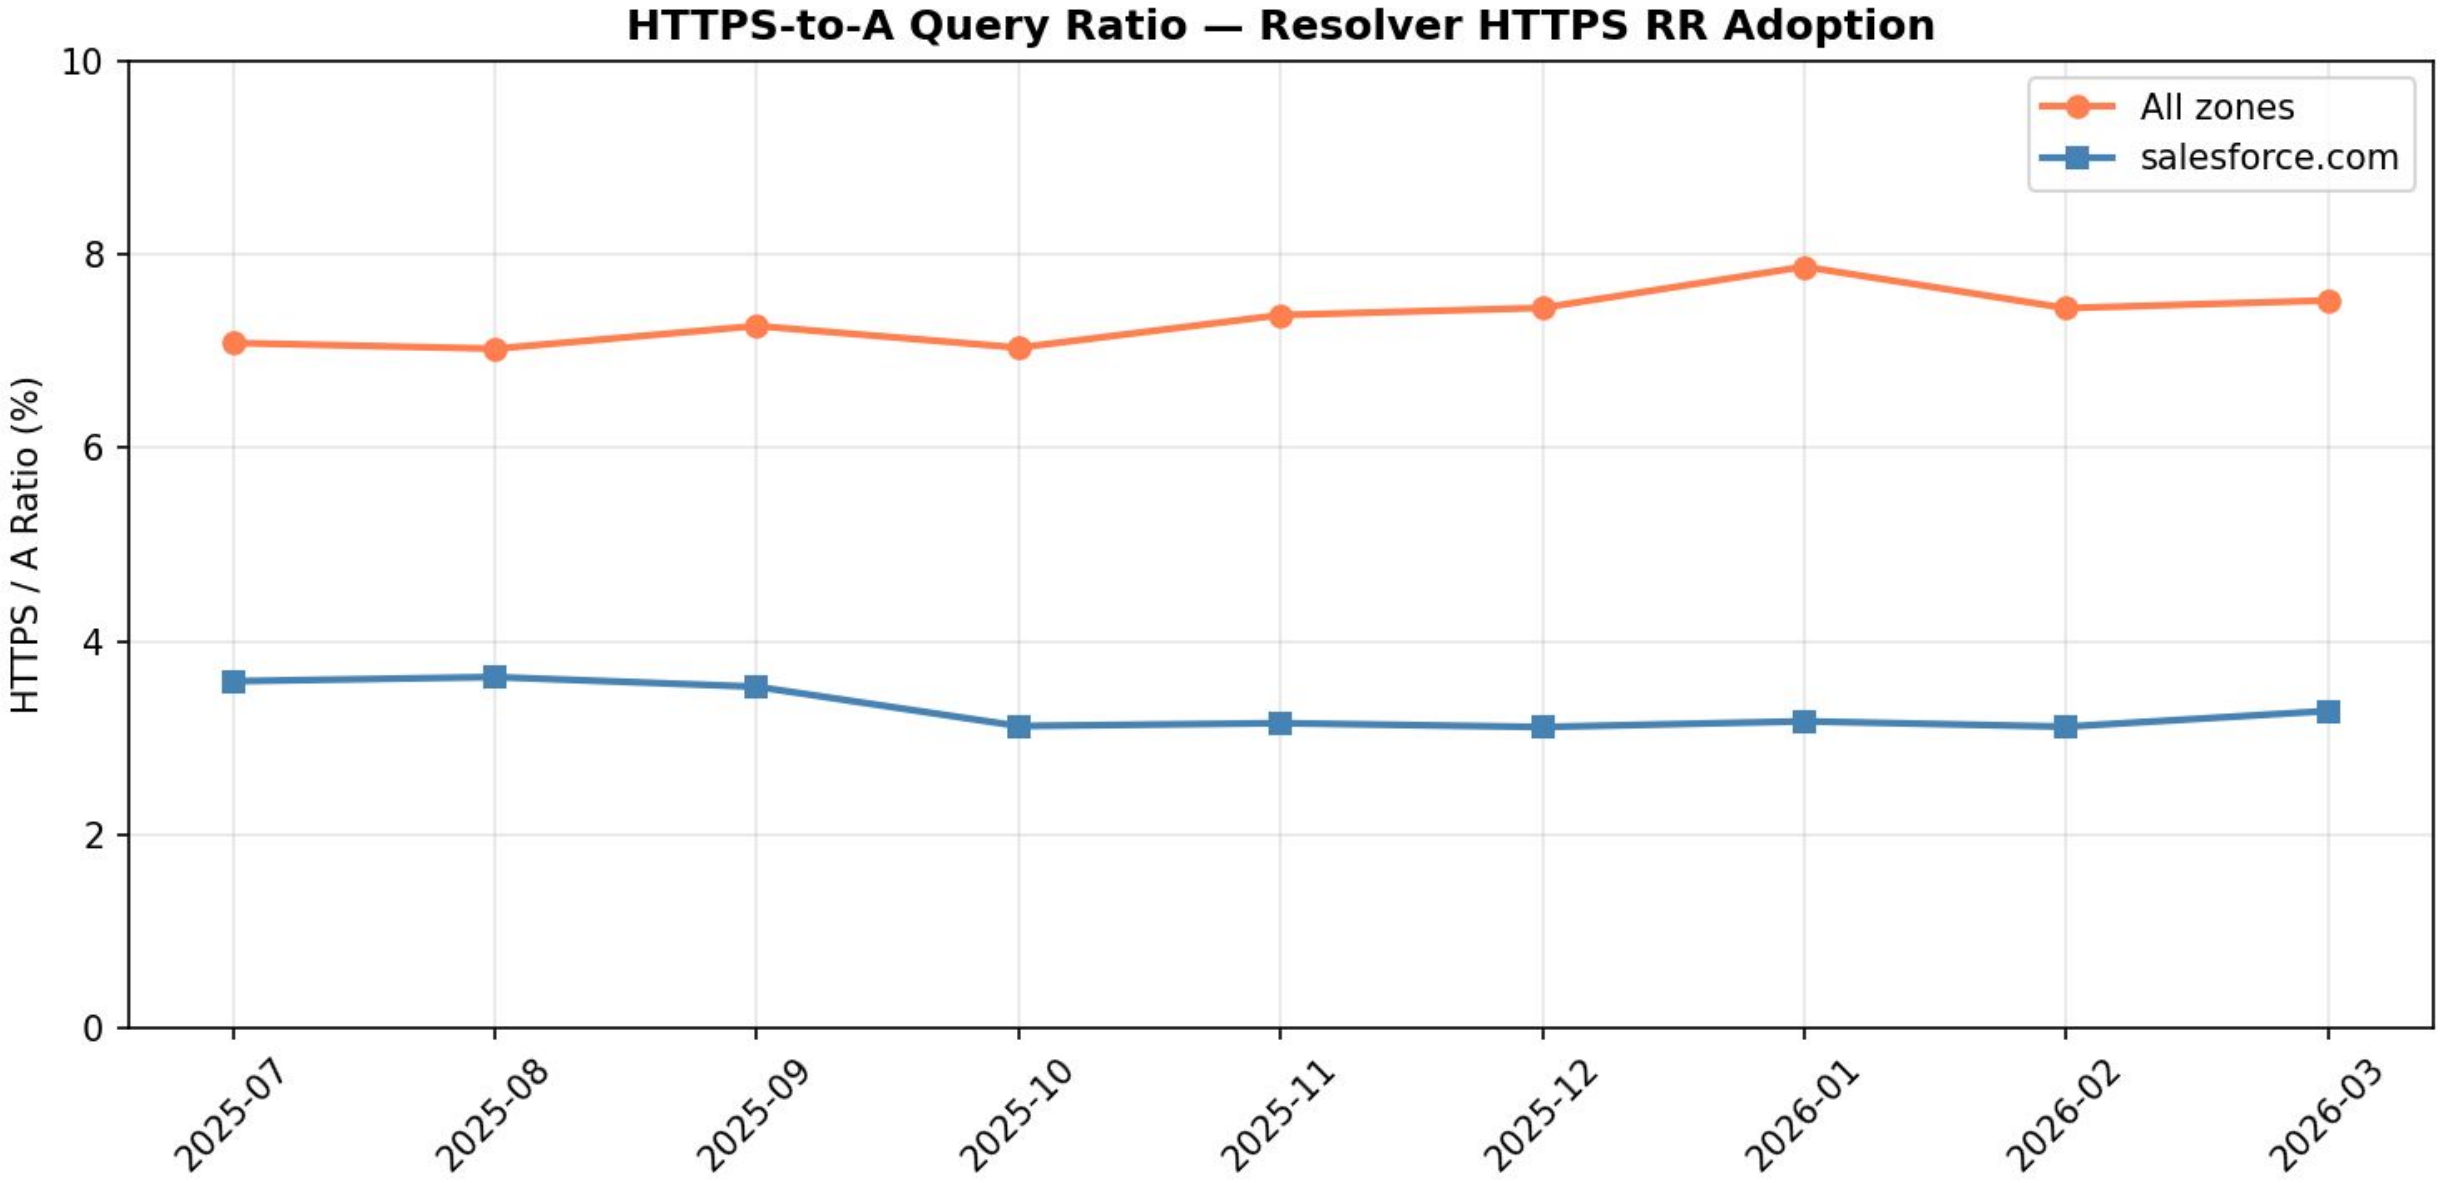
HTTPS-to-A Query Ratio — Resolver HTTPS RR Adoption
HTTPS / A Ratio (%)
10
8
6
4
2
0
All zones
salesforce.com
2025-07
2025-08
2025-09
2025-10
2025-11
2025-12
2026-01
2026-02
2026-03

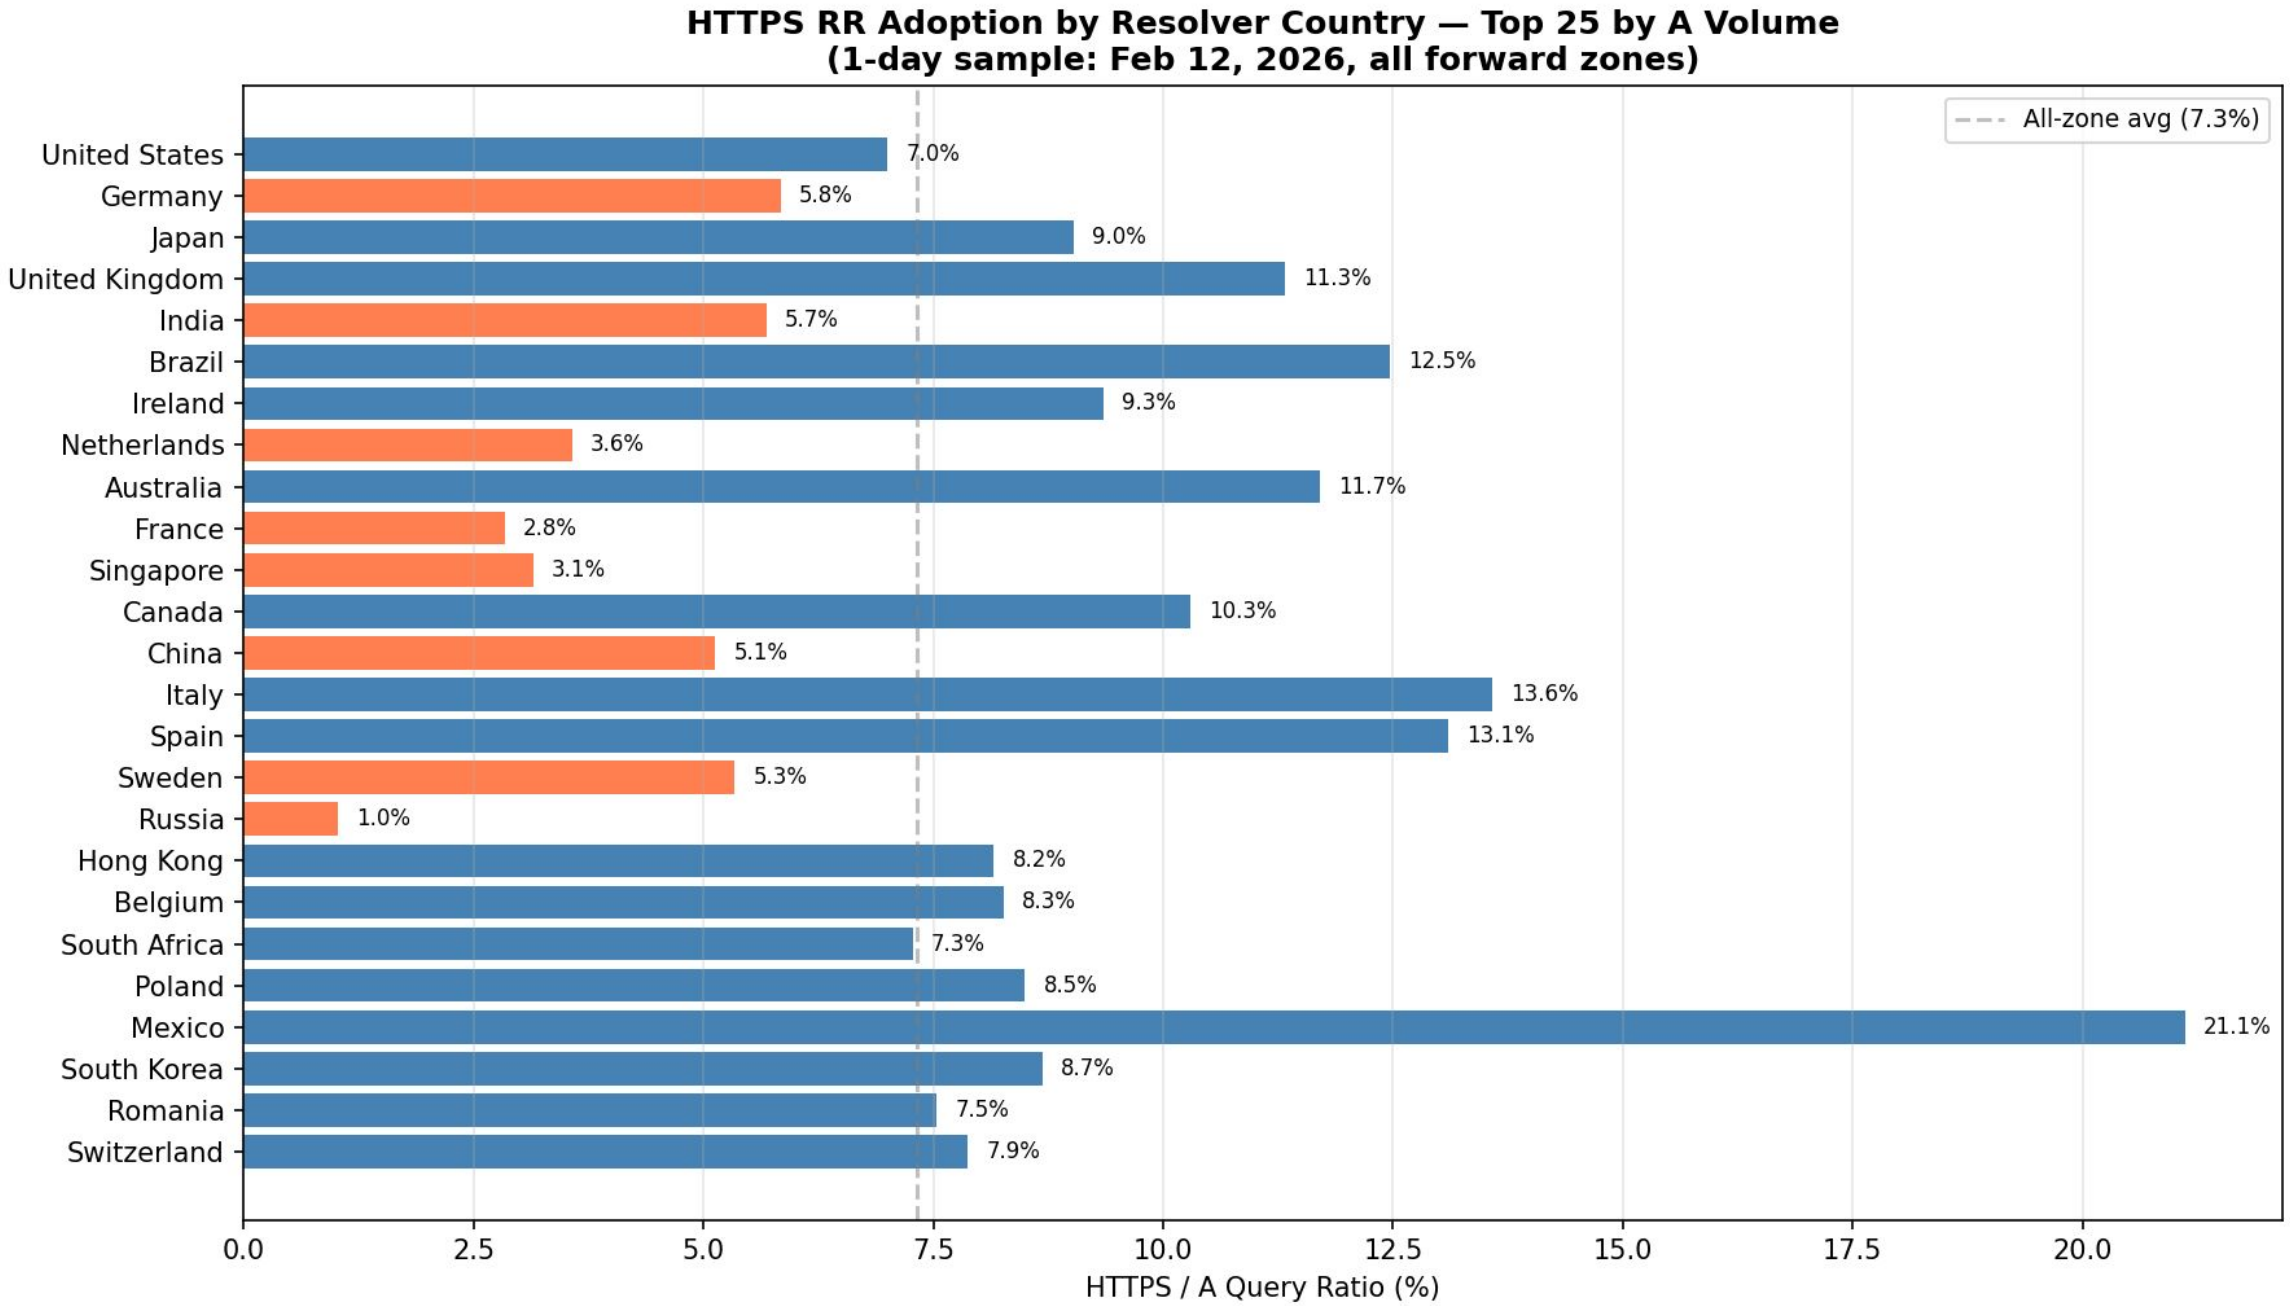
HTTPS RR Adoption by Resolver Country — Top 25 by A Volume
(1-day sample: Feb 12, 2026, all forward zones)
All-zone avg (7.3%)
United States 7.0%
Germany 5.8%
Japan 9.0%
United Kingdom 11.3%
India 5.7%
Brazil 12.5%
Ireland 9.3%
Netherlands 3.6%
Australia 11.7%
France 2.8%
Singapore 3.1%
Canada 10.3%
China 5.1%
Italy 13.6%
Spain 13.1%
Sweden 5.3%
Russia 1.0%
Hong Kong 8.2%
Belgium 8.3%
South Africa 7.3%
Poland 8.5%
Mexico 21.1%
South Korea 8.7%
Romania 7.5%
Switzerland 7.9%
0.0
2.5
5.0
7.5
10.0
12.5
15.0
17.5
20.0
HTTPS / A Query Ratio (%)

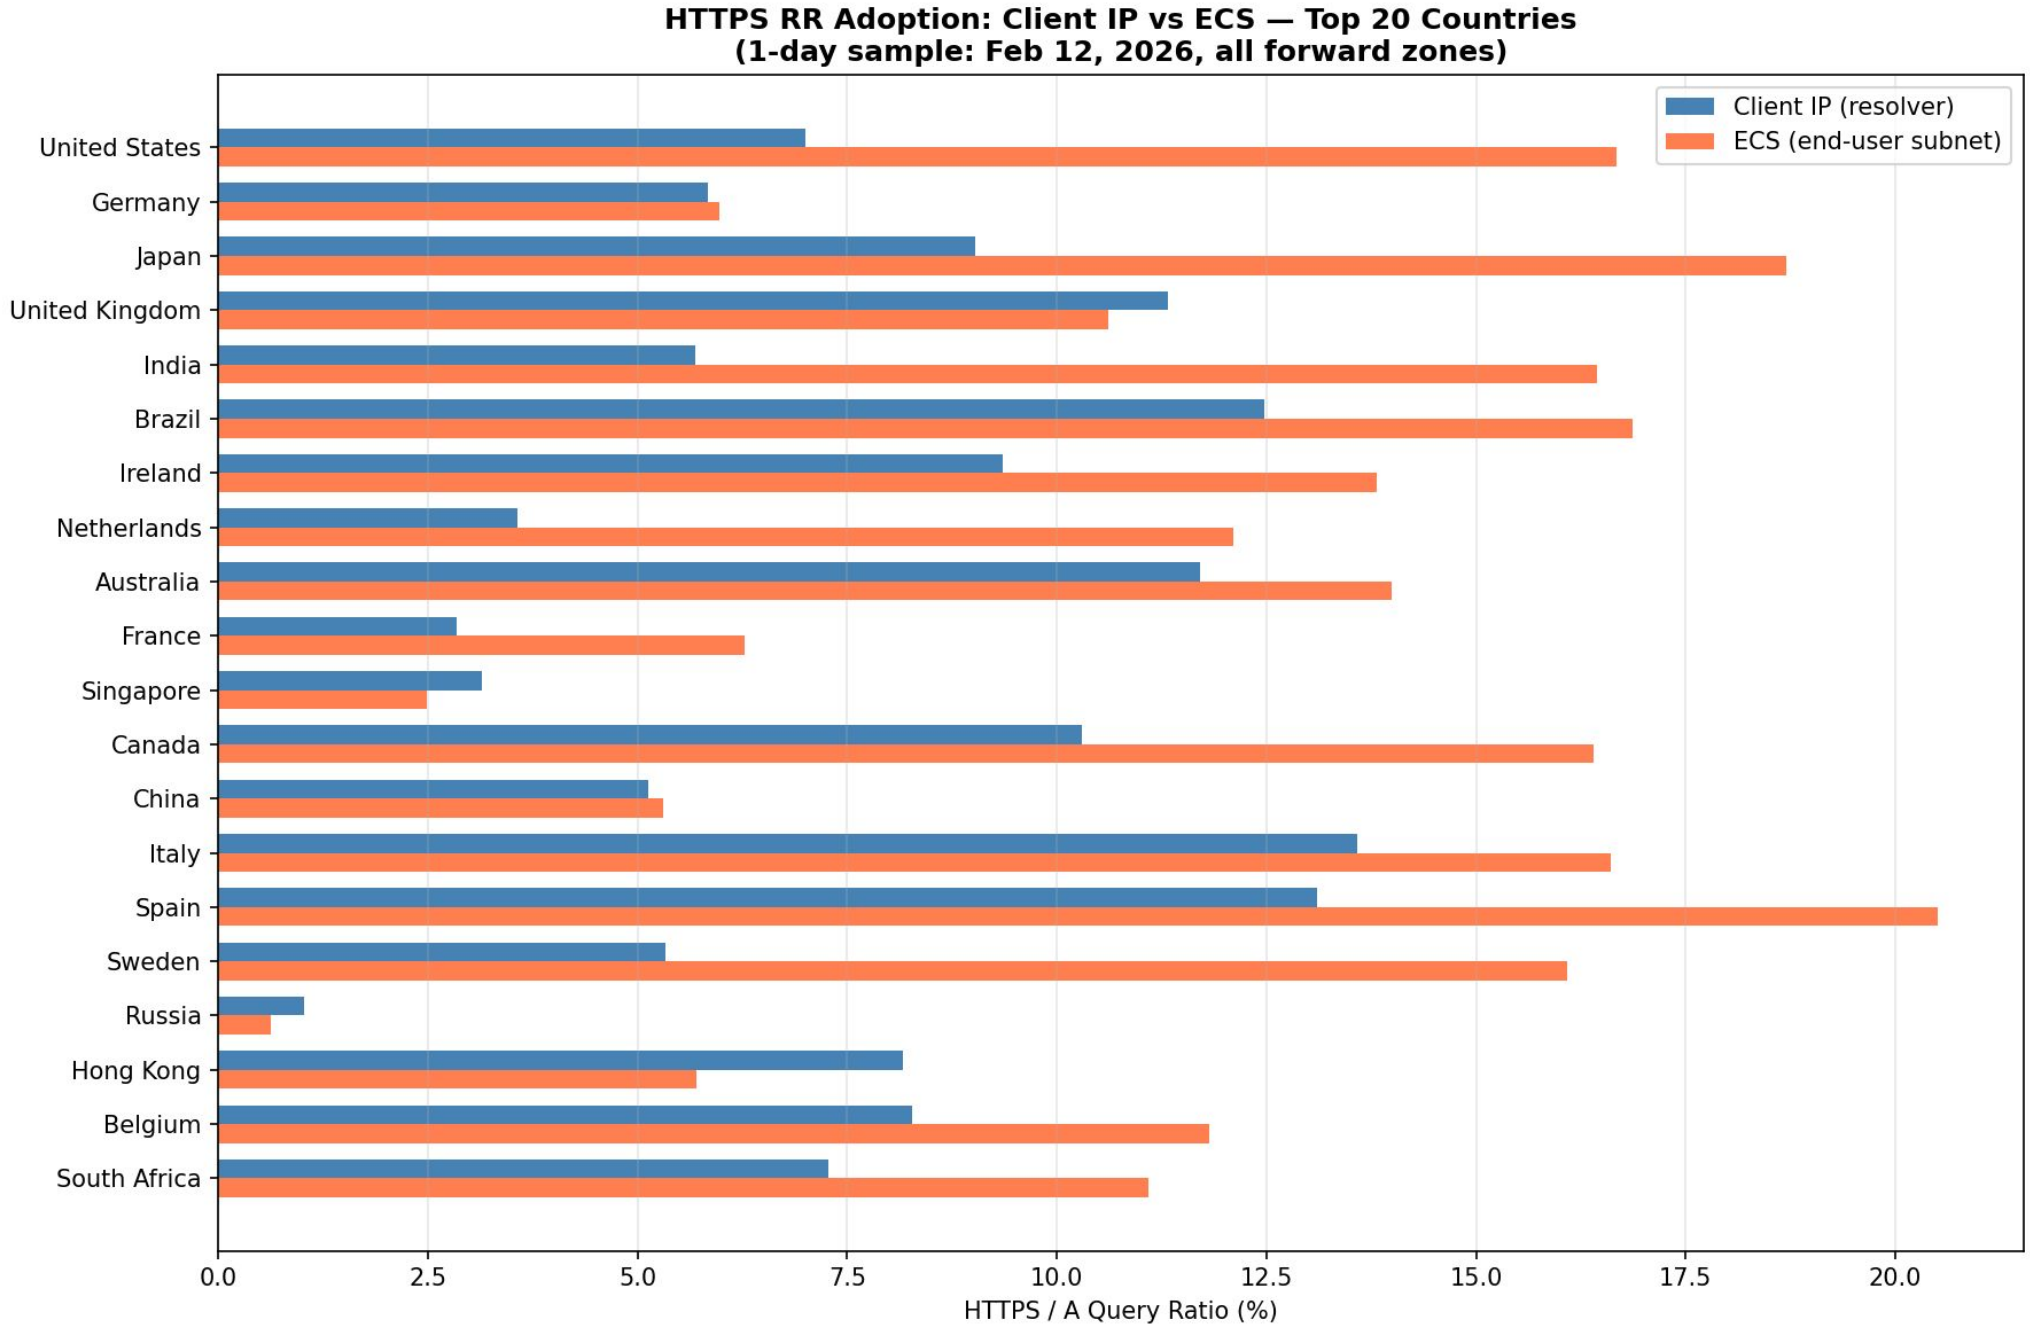
HTTPS RR Adoption: Client IP vs ECS — Top 20 Countries
(1-day sample: Feb 12, 2026, all forward zones)
Client IP (resolver)
ECS (end-user subnet)
United States
Germany
Japan
United Kingdom
India
Brazil
Ireland
Netherlands
Australia
France
Singapore
Canada
China
Italy
Spain
Sweden
Russia
Hong Kong
Belgium
South Africa
0.0
2.5
5.0
7.5
10.0
12.5
15.0
17.5
20.0
HTTPS / A Query Ratio (%)

EDNS Client Subnet Usage

Only 8.6% of observed traffic carry the EDNS client subnet option

94.6% of Resolver IPs never send ECS (so 5.4% send)

ECS volume primary comes from a very small set of public resolvers

Including some that publicly state that they don't (selective, per domain)

Some of these resolvers send it selectively, others send it always.

ECS based traffic profile for DNS is often significantly different from resolver IP

e.g. Observed HTTPS/A ratio is almost double, reflecting a different aggregate type of clients (consumer browser dominance most likely)

Inactive Zones

Inactive zones, or zones with only erroneous or NXDOMAIN queries, and who is generating those queries.
Analysis revealed many of these are obsolete or should otherwise be deleted.

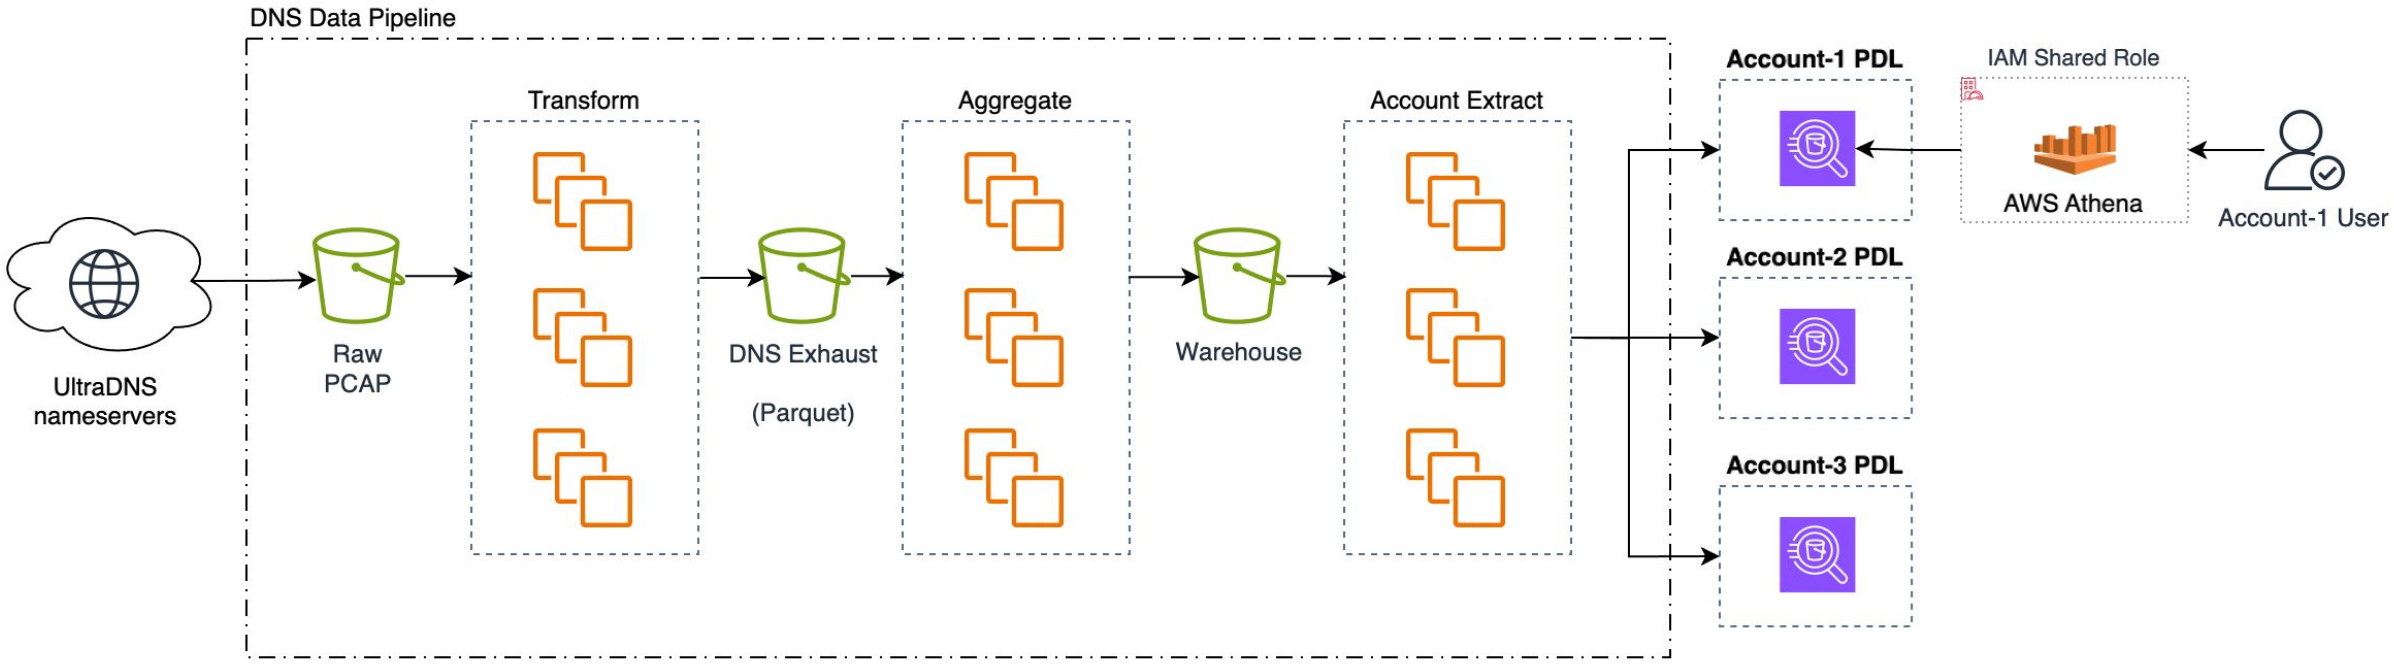

Transform

- Parse packets
- Match queries/responses
- Deduplicate
- Enrich – zone, account, node, geo
- Output parquet format column data

Aggregate

- Minute, Hour, Day
- Node
- Account
- Zone
- Add to exhaust data

Extract

- For single account
- DNS responses from exhaust (parquet)
- Aggregate – minute, hour, day, zone
- Write to account isolated storage

Analyze

- Shared AWS Athena
- 12-month history retention
- Data updated hourly
- DNS responses
- Client Geo and Network tags

Data Collection Architecture

Limitations/Improvements

Some analyses we wanted to perform but were unable to:

- Examining client advertised EDNS buffer size (would improve our truncation analysis)
- Examine range of EDNS header flags and options in use (selective collection of only common ones) - DELEG, CDoE, etc
- Response size.
- Detecting probing for secure transport (DoT, DoQ, etc.)
- Non realtime nature of the data collection
  - Complex ETL pipeline with inherent data lag (a few hours currently)
- In discussions with UltraDNS and extending their data collection & process to address of these limitations.

General comments

Multiple DNS vendors, but only looking at one.

We suspect broad similarities in traffic profiles though.

These are DNS traffic characteristics peculiar to our enterprise DNS infrastructure.

How representative are they of other enterprises/SaaS providers? We don't know.

We do know that we diverge from different types DNS infra (like the root and TLDs for example, where the NXDOMAIN/NODATA distribution is significantly different from what we've shown).

Plan to continue longitudinal measurements and analyses.

Capabilities across other managed DNS vendors

Most vendors offer some form of near realtime query logging capability. But the captured data is fairly limited, and there are often storage/cost challenges doing this at scale.

Furthermore none today offer a "data lake" type of capability, with detailed additional info. Per query/response data, raw packet and pre-aggregated tables, metadata enrichment (geo/carrier info etc), sql queriability, turnkey configuration, centralized multi-account data collection, etc.

General Capability across the industry?

Should other DNS platform vendors have a Private Data Lake type feature? We think so (probably as a premium product).

Large well resourced orgs have the resources to utilize this effectively, and in this age of A.I. coding agent tools, this kind of analysis is well within the reach of almost any organization.

What is the data analysis good for?

- Traffic identification, trends, characterization
- Detecting misconfiguration in DNS zone data and/or applications
- Security analysis and response
- Informing applications about feature adoption
- Research
- Marketing?

Acknowledgements

Steve De Jong & Chester Hockersmith of Digicert/UltraDNS.

A Look at Traffic to Authoritative DNS Servers of a Large Enterprise

Thank you!

Questions or Comments?