

BigTime: Characterizing Large Time Service Providers

Technical Report SIDN Labs 2026-06-23 Updated from 2025-12-01

PASCAL HUPPERT, University of Twente and University of Münster, The Netherlands / Germany

GIOVANE C. M. MOURA, SIDN Labs and TU Delft, The Netherlands

MARCO DAVIDS, SIDN Labs, The Netherlands

PIETER-TJERK DE BOER, University of Twente, The Netherlands

RALPH HOLZ, University of Münster and University of Twente, Germany / The Netherlands

REIN FERNHOUT, University of Twente, The Netherlands

GEORGIOS SMARAGDAKIS, TU Delft, The Netherlands

CRISTIAN HESSELMAN, SIDN Labs and University of Twente, The Netherlands

The Network Time Protocol (NTP) keeps Internet-connected clocks synchronized, which core applications and protocols depend on for proper functioning, including DNS caches, RPKI, and TLS. Next to free time servers (NIST, NTP Pool), large cloud providers and OS vendors such as Microsoft, Apple, and Google have begun operating their own time services and setting their billions of systems/devices to use their NTP servers by default. This creates a level of service centralization around these providers, and yet there has been little scrutiny of these services and their characteristics. In this paper, we characterize seven large time service providers in terms of replication architecture, client mapping, accuracy and service provided. We found a large variation in service replication (DNS and anycast), how they map clients to servers, and in their accuracy and (security) feature support. This includes violations of best practices, servers clocks being incorrect by up to 50ms on one occasion, the use of outdated NTP versions, and a lack of support for NTS, RPKI, and DNSSEC.

1 Introduction

Synchronized clocks are essential for modern societies, being required for financial transactions, power grid operations, and telecommunication networks [29]. On the Internet, core services and applications depend on clock synchronization to verify validity [20, 40, 49, 77], such as TLS [71], DNSSEC signatures [4], DNS caches [58] and RPKI [12]. Various services and applications can fail when wrong time information is served to clients, as in 2012 at the US Naval Observatory (USNO), when their time services published wrong time information (off by 10 years), causing many network routers and Active Directory servers to experience outages [6, 44].

The Network Time Protocol (NTP) [49] is the Internet's default protocol for clock synchronization. NTP clients request time information from NTP servers, which, in turn, are synchronized with out-of-band high precision references, such as atomic clocks or GNSS like GPS and Galileo [49]. There have been many public NTP services on the Internet for decades, including NIST [60] in the US, NPL [37] in the UK, PTB [10] in Germany and the NTP Pool [63].

More recently, device vendors such as Apple, Microsoft, Google, and Ubuntu started to provide their own time services [55] and, more importantly for our study, *setting them as default in their devices*, sometimes even hardcoding them – for instance, Apple's iOS cannot be reconfigured [38]. Other cloud and content providers have also started their own time services, including Cloudflare [16], AWS [78], and Meta [65].

Authors' Contact Information: Pascal Huppert, University of Twente and University of Münster, The Netherlands / Germany, pascal.huppert@utwente.nl; Giovane C. M. Moura, SIDN Labs and TU Delft, The Netherlands, giovane.moura@sidn.nl; Marco Davids, SIDN Labs, The Netherlands, marco.davids@sidn.nl; Pieter-Tjerk de Boer, University of Twente, The Netherlands, p.t.deboer@utwente.nl; Ralph Holz, University of Münster and University of Twente, Germany / The Netherlands, ralph.holz@uni-muenster.de; Rein Fernhout, University of Twente, The Netherlands, r.p.j.fernhout@student.utwente.nl; Georgios Smaragdakis, TU Delft, The Netherlands, g.smaragdakis@tudelft.nl; Cristian Hesselman, SIDN Labs and University of Twente, The Netherlands, cristian.hesselman@sidn.nl.

Provider	Domain Name	User Base
Microsoft	time.windows.com	1.4B [45]
Apple	time,time-[macos,euro,ios].apple.com	2.2B [3]
Google	time.android.com,time.google.com	3.0B [8]
Ubuntu	ntp.ubuntu.com	Unclear
AWS	time.aws.com	-
Cloudflare	time.cloudflare.com	-
Meta	time,time[1-5].facebook.com	-

Table 1. Evaluated time service providers. Providers highlighted are enabled by default in their respective operating systems (OSes).

The implications of these settings are significant: billions of devices are configured to *use only* their vendor-provided time servers, regardless of network or geographic location. As shown in Table 1, there are at least 6.6B active devices from Apple (laptops, smartphones and watches), Microsoft (servers, laptops, and smart devices), and Google (Android) configured by default to use their vendor’s time servers.

Compared to DNS, where clients typically change resolvers when moving across networks, NTP settings mostly remain constant (although DHCP [1] allows operators to configure provide NTP servers, clients may simply ignore it). Therefore, NTP centralization is *far more pronounced* than in other services, such as the DNS [54]: a small number of vendors effectively serve as the time authorities for billions of devices (Table 1). This makes the study of these time services even more pressing.

While a previous study [55] has examined the NTP Pool, which is widely used among Linux distributions, to date there have been no studies scrutinizing the time services provided by these other vendors and cloud/content providers – how their time services are set up, how they map clients to servers, and how accurate they are.

Therefore, in this paper we characterize the time service networks of large commercial vendors and cloud operators – “providers” hereafter (Table 1). We investigate their architecture – how they replicate their services to serve their billions of clients (§3) – and compare this to a well-documented replicated service: the Root DNS system [76]. We investigate the criteria used by each time service provider to map clients to specific NTP servers (§4). Finally, we evaluate the accuracy provided by each time provider and their services characteristics (§5).

We make the following contributions:

- We show a large diversity in terms of replication: out of seven providers, three use IP anycast [66], whereas four use multiple (up to 90) IP addresses. For the unicast-based providers, we find that client geolocation is the primary method to map clients to server, just as for the NTP Pool.
- We show how Microsoft’s time service served half of our 9k vantage points (VPs) with a single NTP server IP address, providing no redundancy and violating NTP best practices [70].
- We measure the accuracy of the providers on four GPS-synced VPs and show they deliver accurate time. However, we found that three of Microsoft’s servers were out-of-sync for several hours during one measurement, by up to 50ms.
- We show the diversity of NTP services provided in terms of stratum of the servers, NTP version, and IPv6 support. Only Cloudflare deploys the latest standards: Network Time Security (NTS) [22], RPKI, DNSSEC, and IPv6. Microsoft is the only provider not supporting IPv6, and provides NTP version 3 only (all others provide NTPv4).

2 Background

NTP is designed to synchronize the clocks of computer systems over variable-latency networks [49]. It provides an accuracy of milliseconds from Coordinated Universal Time (UTC). It uses a hierarchical system of time sources and servers ordered into levels by *stratum*, which indicates the distance from the reference clock. Stratum 1 servers are directly connected to a time source, for instance an appliance that uses Global Navigation Satellite Systems (GNSS) like GPS and Galileo. Stratum 2 servers are synchronized with stratum 1 servers.

Client-Server synchronization: NTP clients synchronize their clocks by exchanging timestamped requests and responses with NTP servers. Each exchange yields four timestamps: the origin timestamp T_1 (client time when the request is sent), the receive timestamp T_2 (server time when the request arrives), the transmit timestamp T_3 (server time when the response is sent), and the destination timestamp T_4 (client time when the response arrives), as can be seen in Figure 1. Assuming symmetric network delay, the client computes the offset between its clock and the server's as θ and may adjust its clock by θ to align with the server.

Protocols: The Simple Network Time Protocol (SNTP) [48] is a simplified version of the Network Time Protocol (NTP) designed for less powerful computers and applications where accuracy of clock synchronization is less critical. SNTP uses the same packet format and general process as NTP, and differs mainly in its client implementation. A recent non-peer-reviewed study [35] has shown how clients of all major OSes (Windows Server, macOS and Ubuntu up to Oct. 2025) use only SNTP-based clients which are vulnerable to time-shift attacks.

Network Time Security (NTS) [22] is an extension to NTP designed to provide cryptographic security for time synchronization. It uses TLS to ensure authenticity, integrity and partial confidentiality of time-stamped messages exchanged between clients and servers, protecting against various attacks such as replays and man-in-the-middle attacks. Of the time service providers from Table 1, only Cloudflare and Ubuntu currently support NTS. (Ubuntu has started support in Oct. 2025 [43]).




Fig. 1. Timestamps used in NTP offset calculations

3 Time Services Architecture

The commercial time services we evaluate (Table 1) serve billions of devices daily and are therefore likely highly replicated. In this section, we carry out measurements to determine their replication choices. Before we do that, we summarize the replication strategies used by two well-known services: the NTP Pool [63] and the Root DNS server system [76].

3.1 Examples of highly replicated services

NTP Pool: The NTP Pool uses DNS-level replication to serve its clients. It has 3,800 IPv4 NTP servers (as of 2026-06-06), all associated with the `pool.ntp.org` domain name. Clients are mapped to a subset of these servers based on their geographical location [55]. Therefore, the NTP Pool uses DNS-level replication, associating multiple IP addresses with the zone.




Fig. 2. Service replication in the Root DNS system.

The Root DNS system: The Root DNS system [76] comprises 13 IP addresses with 2,015 servers (as of 2026-06-06), referred to as instances. The root servers deploy three levels of replication (Figure 2): DNS-level, IP anycast [66], and site-level (local load balancing).

At the DNS-level, the system's only domain name (the dot `.`) is associated with 13 IPv4 and 13 IPv6 addresses. Each of these IP addresses is further replicated with IP anycast, which allows the same routing prefix to be announced from multiple locations on the globe, with each location referred to as an anycast site (site1–site3 in Figure 2). For instance, F-Root [17] has 366 anycast sites as of 2026-06-06. Each of these anycast sites may employ further local replication using load balancers (s1–s3 in Figure 2). For example, K-Root [14] uses three instances for its Amsterdam site.

3.2 DNS and Anycast Replication

We evaluate how the time service providers (Table 1) compare to these well-studied examples.

3.2.1 DNS-level replication. We first evaluate if the provider deploy DNS-level replication. To determine whether they rely on the client's geographical location when mapping users to servers as the NTP pool does, we configure 500 RIPE Atlas probes [74, 75] as vantage points to send one DNS query per 10 minutes, for 8h (datasets: [72]). We configure each vantage point to query one of the authoritative DNS servers of the time service directly¹. This way, we bypass the recursive resolvers' caches and retrieve fresh responses with each new query.

For each provider, we then compute the number of unique IP addresses they return over the course of the measurement. Table 2 summarizes the results. We notice clear differences: AWS has 90 IPv4 servers, followed by Apple (53) and Microsoft (12). Google, despite likely having the largest client population, has only 4 IPv4 address.

3.2.2 Anycast replication. To determine if the IP addresses obtained from our Atlas measurements are deployed using anycast §3.2, we check them against IPInfo [32], the BGPTools anycast prefix list [5], and the LACeS Anycast census [26, 80]. We consider an IP as anycast if it is listed as such in one or more of those data sources.

Table 3 summarizes the results. We see that Google, Cloudflare, and Meta deploy IP anycast for their time services. For each network prefix, we report the maximum number of anycast sites as reported in March 2025 by the LACeS data [81], which we show in Table 3. For instance, we see that Google has 41 sites, far more than IP addresses. Note that these are lower bounds, as measuring anycast deployments requires a large number of vantage points and depends on BGP routing. For

¹For Microsoft and Apple, we do not configure the VPs to query the domains shown in Table 1 but the *actual* domains used. Microsoft's `time.windows.com` domain name redirects to `twc.trafficmanager.net` using a CNAME record. Apple's domains redirects to two domains: `[time, time-osx].g.aaplimg.com`. However, we find that both domain names point to the same NTP servers. Hence, we just report on `time.g.aaplimg.com` and include `time-osx.g.aaplimg.com` in Appendix A and Appendix B

	Domains	IPv4	IPv6	ASes v4	ASes v6
Microsoft	1	12	0	1	0
Apple	4	53	48	2	2
Google	2	4	4	1	1
Ubuntu	1	4	3	1	1
Amazon	1	90	90	2	2
Cloudflare	1	2	2	1	1
Meta	1	5	5	1	1

Table 2. DNS-level replication of time service providers.

Provider	Prefixes	Anycast	Sites (IPv4)	Sites (IPv6)
Microsoft	11/0	No	–	–
Apple	53/48	No	–	–
Google	1/1	Yes	41	No data
Ubuntu	4/3	No	–	–
AWS	46/9	No	–	–
Cloudflare	1/1	Yes	63	47
Meta	5/5	Yes	8	11

Table 3. Anycast replication of time service providers.

	Microsoft	Apple	AWS
Domain	twc. traffic-manager.net	time.g. aaplimg.com	time.aws.com
VPs	9,010	8,994	8,939
Countries	170	170	171
ASes	3,094	3,094	3,081
DNS Queries	642,617	642,042	640,253
NTP Servers	12	53	90
ASes	1 (AS8075)	2 (AS6185 & AS714)	2 (AS14618 & AS16509)
Dataset	Microsoft-map	Apple-map-1	AWS-map
Auth server	13.107.222.240	17.253.206.8	205.251.193.84

Table 4. Atlas measurements. Datasets: [72]

Meta, the small number of sites is likely accurate, but for larger numbers, the anycast census might report significantly smaller numbers, meaning that Google and Cloudflare could have multiple hundreds of sites [27]: Statistics about Google’s network [24] suggest the actual number may be up to 200 sites, while Cloudflare claims to have 330 locations [16].

Unicast services: The remaining providers (Microsoft, Apple, AWS, and Ubuntu) all use IP unicast services. Among these, Ubuntu uses only 4 IP addresses for the entire client population, whereas Microsoft and Apple, with more than 1B clients each, have 12 and 53, respectively. Given Apple and Microsoft serve billions of devices, we expect them to have locations worldwide. We show the geographical location of their services and AWS in Figure 3 and for Ubuntu in Appendix A. We see Microsoft has no servers in the southern hemisphere, and Microsoft is the only provider which does not support IPv6.




Fig. 3. Unicast NTP servers geolocation.

Verifying unicast providers size: To determine the number of servers for the unicast providers (Microsoft, Apple, and AWS, [Table 2](#)), we carry out additional measurements from a larger set of VPs: 9k Atlas Probes from 170 countries and 3k ASes for each time provider. [Table 4](#) summarizes the results. We send 640k DNS queries per provider. After data sanitation (excluding IP addresses that do not belong to the time service provider’s own Autonomous System (AS) and likely are received due to interception of client DNS queries [[56](#), [61](#)]), we identify that Microsoft, Apple, and AWS return 12, 53, and 90 IPv4 addresses, [Appendix B](#) shows how often this happens for each NTP server.

While this measurement spans an 8-hour window; a natural concern is whether this is long enough to observe the full set of NTP servers behind each service. To test this, we ran an additional lower-frequency campaign over two months across April and May 2026, using the same 8.9k Atlas probes but issuing queries three times per day rather than every ten minutes (datasets: [[73](#)]). Even over this longer period, we observed similar server counts—12 for Microsoft, 52 for Apple, and 86 for AWS—confirming that an 8-hour window already captures the complete set. We report the full details in [Appendix D](#).

4 Client/Server Mapping


Next we set out to determine how each time service maps clients to their available servers. As a comparison, the NTP Pool maps based on their geolocation: clients are mapped either to servers in the same country, or the same continent in case of no servers being run in the country [[55](#)]. (The Pool operators state they employ this criteria to avoid risks of packet loss and issues that asymmetric routing can cause [[25](#)]).

4.1 Mapping Criteria

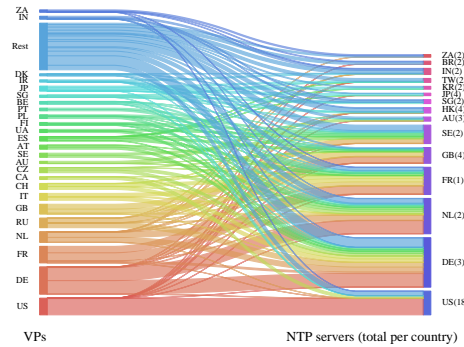
We investigate the criteria used by Microsoft, Apple, and AWS to map their clients to their NTP servers ([Figure 3](#)). We use the Atlas measurements from [Table 4](#). For each Atlas VP, we compare its geographical location (from probe metadata [[15](#)]) to the geolocation of their designated NTP servers. (We do not evaluate Anycast providers given the mapping is done by BGP; and Ubuntu, given it always returns all four NTP server addresses).

[Figure 4](#) shows the results. The left side of the Sankey diagrams contains the countries of the Atlas VPs (which represent real clients) and the right side shows the NTP servers’ countries. Analyzing these figures, we can see that these providers seem to use the client’s geolocation to assign them to specific NTP servers.

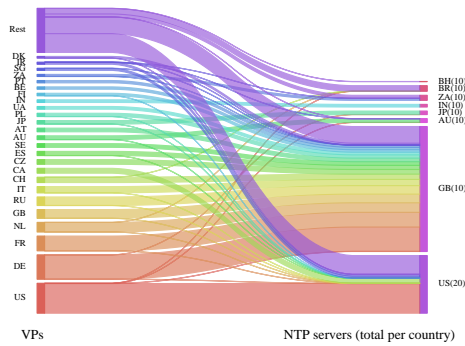
In-country mappings: For clients located in countries where the time providers operate a server, we see that, just like the NTP Pool, they are mapped to in-country servers. For example, Microsoft maps 1,297 out of 1,313 US-based VPs (98.7%) to US-based servers, AWS mapping 1,305 out of 1,311



(a) Microsoft. Dataset: Microsoft-map [72]



(b) Apple. Dataset: Apple-map-1 [72]





(c) AWS. Dataset: AWS-map-1 [72]

Fig. 4. Client-server mapping: sankey diagrams of Atlas Probe VPs and NTP Servers. We show only countries with more than 70 VPs, the remaining grouped as Rest

(99.5%). Apple, in turn, seems to do the same for the US-based VPs, but not for Europe: European VPs are distributed among servers in multiple countries in the region – for instances, we see 1,036 German VPs being mapped to both Germany (DE) and the Netherlands (NL). Apple seems to group all of Western Europe together in these mappings (Figure 4b). Considering that the area of Western Europe is 1/8 of the continental US area, this seems reasonable geographically.

Out-of-country mappings: For clients without NTP servers in their country, we see a more diverse mapping. Microsoft behaves similar to the NTP Pool, mapping clients to their continent in case of no in-country servers. We see that German VPs are mapped mostly to servers in the Netherlands. Russia, which spans Europe and Asia, is mapped to both continents. As we showed in Figure 3a, Microsoft has no servers in South America and Africa. Most of African VPs are mapped to Europe, while most of South American VPs are mapped to the US servers. It is worth noting that Microsoft’s time service domains point to a domain name associated with Azure’s DNS-based traffic distribution (`twc.trafficmanager.net`), which supports geolocation-based mapping [46].

Multi-continent mappings: Apple maps clients from EU countries with no NTP servers to mostly EU countries (ES VPs are mapped to FR, GB, and DE). It has NTP servers in Africa (South Africa), but maps African VPs to virtually all continents: South African VPs are mapped both to South



(a) Number of NTP servers Atlas VPs are served (measurement from Table 4). (b) Geolocation of Atlas VPs served only a single time server from Microsoft’s time services.

Africa (91), Brazil (53), and Asia (38). South American VPs are mostly mapped to both Brazil (same continent) and North America. Iranian VPs (77), are mapped to both India and Sweden – so two different continents. This contrasts with the NTP Pool, which map all clients from a specific country to the same single continent, in case there are no server in the client’s country.

For AWS, most of EU-country probes are mapped to servers in Great Britain (GB) – the only European NTP server location for AWS (Figure 3c), despite AWS having presence in other EU countries. South American VPs are mapped to Brazil’s location (23 from Argentina, 59 from Chile), and African VPs are mapped to South Africa’s servers. AWS differs from the NTP Pool in regards to Asia (and the Middle East): despite having server in Bahrain, India, and Japan, it maps many Asian countries to North America, including Singapore, Indonesia, Iran and China. This also differs from the NTP Pool’s mapping.

Takeway: Amazon and Apple, the unicast-based service providers with dozens of sites, use the client’s geographical location to map them to “nearby” servers, similar to the NTP Pool. We see some improvements that differ from the NTP Pool (handling Europe as a continent instead of individual countries and mapping clients to multiple continents). However, it is interesting to observe that both the NTP Pool and the large commercial providers in essence perform a similar mapping: the client’s geographical location is the primary criteria.

4.2 Number of servers per client

BCP223/RFC8633 (§7 in [70]) discusses the NTP best current practices and expressly states that each client should use multiple time servers. We evaluate whether providers enable this by serving all clients with more than one time server.

All anycast providers return multiple IP addresses (Table 2) to each client, and so does Ubuntu (unicast). Therefore, only Microsoft, AWS, and Apple are left to be assessed, which have from 12 to 90 NTP server addresses (IPv4).

We analyze the datasets from Table 4, which we drew from more than 9k Atlas VPs, to determine how many servers their clients are served, from all the available servers. For each VP, we compute a list of all NTP server’s IP addresses it has been served in the DNS response (each probe sent 72 queries over a 12-hour period, one each 10 minutes, bypassing caching resolvers by querying the authoritative servers directly).

Figure 5a shows the results for these providers. AWS serves each client with 10 addresses, whereas Apple serves most with five addresses. Microsoft, however, serves half of our VPs with a single time source, violating RFC8633. This reduces reliability, given those clients will rely on a single source of time. Considering that Windows has 1.4B active installations, the possibility of so many installations depending on just one source is quite alarming.




Fig. 5. Number of responses containing each NTP server address for each time provider.

To determine if these VPs are bound to specific regions, we show in [Figure 5b](#) the geolocation of VPs served by a single IP address. We see that they are not constrained to any specific regions. We notified Microsoft on 2025-05-02 about the issue and have not heard back.

4.3 Load balancing clients

Given the large levels of replication of the providers, how do they balance client numbers among their available servers? Anycast providers leave this decision to BGP, whereas unicast providers can facilitate the DNS for this. We next evaluate the unicast providers to analyze how they load balance clients.

For each provider, we compute the number of times each server was included in a DNS response, for the datasets of [Table 4](#), which we show in [Figure 5](#).

Microsoft: [Figure 5a](#) shows the client distribution for Microsoft’s servers. We show the NTP server’s geolocation above each bar. We see that the first two servers are returned in the same frequency – both located in the Netherlands. Even if Microsoft has two other servers in Europe (in Great Britain), the Netherlands-based ones are returned more often – German probes could also have been mapped to GB NTP servers, for instance. With regard to the US, we see that one server is returned more often than the other. In conclusion, the client distribution per server demonstrates that Microsoft prioritizes some servers in relation to others, using DNS for load balancing, at least for our set of VPs.

Apple: Apple operates 53 unicast NTP servers. We show the distribution of all of them in [Figure 5b](#) (we include figures for the other Apple domain in [Appendix A](#)). First, we see that there is a clear pattern in load distribution. Given that most US-based servers are assigned to US VPs ([Figure 4b](#)), we analyze how often the servers are returned to these VPs. We see a clear pattern in the frequency of NTP servers: some appear 4× more often than others (48k vs 12k). We see a similar ratio for EU-based servers (155k vs 43k, figure in [Appendix A](#)). We look at the stratum of these servers and see that 4 US-based servers are stratum 2 (shown in red in [Figure 5b](#)), and they are among the ones returned less often.

AWS: AWS has the most clear-cut distribution ([Figure 5c](#)). We see a step-like graph, where each level has 10 servers (we show country codes of only 5 for each group, for readability), and they are returned at the same rates. So clients are redirected to specific locations using the DNS, but served by all NTP servers in that location. Interestingly, despite AWS having two US locations (Virginia and Oregon), the servers in Virginia appear 90× more often than Oregon-based ones.

Takeaway: Load balancing varies among providers, depending on their server locations and number of servers. For our 9k VP set, we see preference towards some servers over others. Anycast-based providers have to rely on BGP to load balance traffic for them.

5 Accuracy and NTP features

Now that we understand the architecture of these time services, we turn to the question of the quality of their service: *How accurate are they?*

Whereas the previous sections relied on DNS measurements, in this section we carry out NTP measurements from multiple VPs synchronized with reference sources (§5.1). We compare the responses to our reference clock to determine their accuracy (§5.2), detect cases of out-of-sync servers (§5.3) and explore the characteristics of services they provide (§5.5).

5.1 Experiment setup

To evaluate the accuracy of time services providers, we have to use VPs which are directly connected to reference sources, having “ground truth” time information. Therefore, we cannot rely on Atlas probes as vantage points, given they have various different methods of updating their clocks [28] but are not directly connected to reference sources.

Therefore, we set up our own VPs, as shown in Table 5, which are connected directly to reference sources, and run Ubuntu Linux and chrony [18] as a time client on both, which is a state-of-the-art NTP implementation [35].

VP locations: Our VPs rely on sources that employ both GNSS and atomic clocks. The first VP is located at AWS Sweden (SE-AWS), and leverages the AWS time service [2]. To prevent bias from the AWS-based location (given AWS is also a time provider), we set up a second VP in the Netherlands (NL-SIDN).

VP	ASN	Time Source	Method
SE-AWS	16509	GNSS, atomic [2]	PHC [2, 69]
NL-SIDN	1140	GNSS, radio, atomic [79]	Linux PTP [69]

Table 5. Accuracy Experiment Vantage Points


Time sources: We connect NL-SIDN to a Multi Reference Source (MRS [33]), which takes time input from multiple time sources, in our case GNSS (GPS and Galileo), Radio (DCF77 [11]), and a rubidium atomic oscillator. Both VPs are synchronized using the hardware-supported Precision Time Protocol (PTP), which is typically at least three orders of magnitude more accurate than NTP [31]. For SE-AWS, we configure chrony to use the PTP Hardware Clock (PHC) [69], whereas for NL-SIDN we configure it with Linux PTP [21].

Measurements: We use a custom Python script to send one NTPv4 query (mode 3, client) to each IP address associated with the time services (Table 2) every 5 minutes for 24h (May 7th 2025 for SE-AWS VP and October 29th 2025 for NL-SIDN VP). In total, each VP sends 288 NTP queries to each of the 322 IP addresses over the 24h period. We capture the traffic using tcpdump and analyze the traces to compute the offset between our VP clocks and the provider’s clocks. We will publicly release the measurement data upon acceptance. We perform another measurement from two more AWS machines (US-AWS in AS14618 and JP-AWS in AS16509, configured like SE-AWS) on June 3rd, 2026, to ensure our results are valid for vantage points in various continents. We remove IP addresses that have become inactive (not responding to NTP queries), meaning these and NL-SIDN have reduced numbers of IP addresses for Amazon and Apple (see Appendix B), which seem to routinely change their server addresses.


Limitations: Our experiment has two limitations. First, our VPs can reach only a single site (location) of anycast providers (Google, Cloudflare, and Meta) each, assuming routing does not change during the measurements [84]. Therefore, we can only partially evaluate these services’ accuracy. (For non-anycast services, we measure all locations by measuring all IP addresses).

Provider	IPv	#IP Addr.	RPKI	DNS-SEC	#Req.	#Resp.	No Resp. #	No Resp. %	Offset mean	Abs. Offset 90th %ile	RTT mean	NTP version	NTS	Stratum
Amazon	v4	89	Yes	No	25 632	25 608	24	0.09%	0.05 ms	1.03 ms	162.05 ms	v4	No	4
	v6	89	Yes	No	25 633	25 613	20	0.08%	0.04 ms	1.04 ms	162.08 ms	v4	No	4
Apple	v4	51	No	No	14 688	14 058	630	4.29%	0.50 ms	6.12 ms	166.54 ms	v4	No	1/2 (78%/22%)
	v6	46	No	No	13 248	12 253	995	7.51%	0.80 ms	5.44 ms	168.75 ms	v4	No	1/2 (77%/23%)
Cloudflare	v4	2	Yes	Yes	576	576	0	0.00%	-0.01 ms	0.06 ms	3.97 ms	v4	Yes	3
	v6	2	Yes	Yes	576	576	0	0.00%	0.01 ms	0.05 ms	3.95 ms	v4	Yes	3
Google	v4	4	Yes	No	1 152	1 152	0	0.00%	0.02 ms	0.79 ms	11.26 ms	v4	No	1
	v6	4	Yes	No	1 152	1 152	0	0.00%	-0.02 ms	0.79 ms	11.12 ms	v4	No	1
Meta	v4	5	Yes	No	1 440	1 440	0	0.00%	0.04 ms	0.90 ms	23.02 ms	v4	No	1
	v6	5	Yes	No	1 440	1 440	0	0.00%	-0.04 ms	1.07 ms	36.24 ms	v4	No	1
Microsoft	v4	12	Yes	No	3 456	3 450	6	0.17%	-13.75 ms	33.25 ms	145.65 ms	v3	No	3
Ubuntu	v4	4	No	No	1 152	1 151	1	0.09%	0.34 ms	3.30 ms	47.45 ms	v4	Yes	2
	v6	3	No	No	864	864	0	0.00%	1.38 ms	1.85 ms	31.01 ms	v4	Yes	2

Table 6. SE-AWS measurement details. Anycasted services are highlighted in orange.



(a) 90th-percentile absolute offset



(b) Average RTT

Fig. 6. Accuracy and RTT results per provider.

Second, while not an inherent limitation of our experiments, NTP assumes symmetric network paths [49]. However, most Internet paths are *asymmetric* [83], often exhibiting different one-way delays ($T_2 - T_1 \neq T_4 - T_3$ in Figure 1). Since the NTP clock offset computation (Figure 1) relies on an equal-delay assumption by averaging the timestamps, persistent delay asymmetries result in systematic estimation errors in the reported offsets [53, 62]. We reduce the impact of this by using multiple vantage points and computing the maximal error in §5.3.

5.2 Accuracy results

Table 6 summarizes our results for the vantage point SE-AWS. We include the same details for the other VPs in Appendix B and discuss the differences in this section.

We evaluate *accuracy* by measuring the clock offset between our VPs and the time service providers (Figure 1). An offset of zero indicates perfect synchronization between the providers' clocks and our own ("ground truth", to within the uncertainty introduced by propagation delays).

In Figure 6a, we show 90th-percentile absolute offset for each provider and both VPs (using absolute values prevents positive and negative values from evening each other out). We see that from each VP, all offsets' 90%iles are under 21ms, which can be considered good for most applications, except for Microsoft from SE-AWS (33.5ms). (We include a figure of the mean offsets in Appendix A, which shows the same pattern).

Accuracy per VP: Comparing VPs, we see that SE-AWS consistently has offsets closer to 0 than NL-SIDN (except for Ubuntu over IPv6), despite having mostly larger average RTTs for the non-anycast




Fig. 7. Violin plots of offset distributions. Positive offset indicates that the polled server’s clock is ahead of our reference clock.

providers (Figure 6b). This might be due to better network conditions at the SE-AWS vantage point, especially for the anycasted providers.

Anycast: For both VPs, we see how anycast providers (Google, Cloudflare and Meta) have a lower RTT, which is due to us only reaching a single site, whereas we measured all locations, globally distributed, for the others (§3.2).

Offsets distribution per provider: Figure 7 shows the offset distribution for SE-AWS and NL-SIDN. We see first a clear difference between the same providers for each VP – NL-SIDN VP has a more dispersed offset for the same providers – including anycast providers, even though it has a lower RTT to many providers.

Microsoft: We zoom in on Microsoft’s servers for both VPs (Figure 7c). We see that for most servers, the NL-SIDN VP has a distribution that is more compact and close to 0, except for two servers (one in JP and other in HK). For our first measurement (SE-AWS), Microsoft’s time servers had predominantly negative offsets, meaning that the server’s clock is lagging behind our reference clock. This will impact the accuracy for clients because Microsoft offers only one server to most (§4.2). The significant differences for Microsoft are likely due to the different measurement times, while the large spread in offsets for two of the servers might be due to large variances in network delays to those specific servers in Asia. We see better accuracy during the follow-up measurement (US-AWS and JP-AWS in Appendix B).

RTT and accuracy: Mathematically, lower RTT improves worst-case accuracy of the offset (Appendix C). This is also the reason behind geo-based mappings in the NTP Pool. Our results, however, show that it depends on the VP. SE-AWS had a weak Pearson correlation between absolute offset and RTT (0.38), whereas NL-SIDN had a strong one (0.75). Therefore, we conclude that this assumption does not always hold, and the VP location also plays a role. This makes intuitive sense: inaccuracies are not caused by RTT, but by asymmetric network latency. We see that despite having 89 IPv4 addresses, 90%ile offset for AWS is under 2ms for SE-AWS (possibly due to bias given both the VP and provider are the same company), but not for NL-SIDN: it delivered the worst offsets (Appendix B).

Stratum and accuracy: From SE-AWS, Cloudflare’s stratum 3 servers provide more accurate results than Google’s stratum 1 (Table 6, Figure 7). This shows that a lower stratum does not always lead to more accurate offsets.

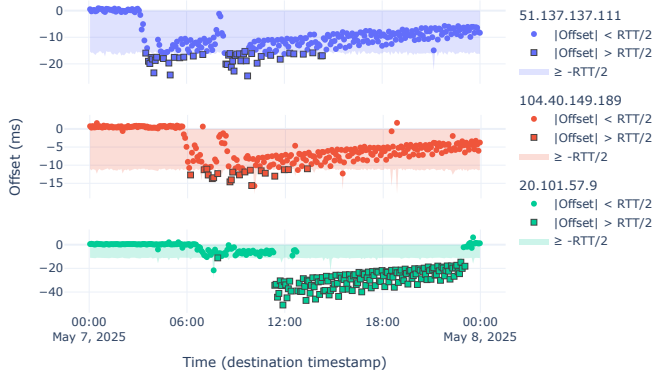


Fig. 8. Offsets of out-of-sync Microsoft servers from SE-AWS. Points with $|\text{Offset}| > \text{RTT}/2$ indicate the server must be out-of-sync.

5.3 Detecting out-of-sync servers

Given that our VPs maintain ground-truth time, any offsets we observe are errors. Such errors can arise from two sources: variable network delay or clock errors on the provider’s side. Because the measured offset combines the server’s true offset with network-delay-induced error, an offset $\neq 0$ does not necessarily indicate that the server’s clock is incorrect. Moreover, without prior clock synchronization, it is generally not possible to distinguish between delay effects and an actual clock error.

However, the contribution of network delay to the offset is limited by the RTT. The offset must fall within $\pm \text{RTT}/2$ (Appendix C, [50]). If the measured offset falls within this delay-dependent bound, then network latency is sufficient to explain the deviation, and the server’s clock may be correct. However, offsets that exceed this bound cannot be explained by latency alone and therefore indicate an unquestionable clock error.

In our measurements, almost all offsets fall within the delay-derived bound, indicating that network latency dominates in the measurement. However, 199 responses from Microsoft servers to SE-AWS violate this, meaning the server’s clocks are out-of-sync. Figure 8 shows the offsets received during our measurement. We provide further statistics about the three affected servers in Appendix B (all of which are in Western Europe, close to the VP). The colored area marks the offset bounds that network latency can cause. Square points are offsets that fall outside the bounds, indicating that the server’s clock is incorrect (the offset cannot be explained by latency alone). The shape of the curve suggests that these servers have synchronized to an incorrect time and slowly approach 0 again.

During the measurements from the other three VPs, we see no further errors from Microsoft. While we also measure some of Cloudflare’s and Meta’s servers to be in error briefly (from JP-AWS and US-AWS, respectively), this never exceeds 2ms ($10\mu\text{s}$ for Meta). Because inaccuracy increases with RTT, we may miss clock errors that are overshadowed by the RTT.

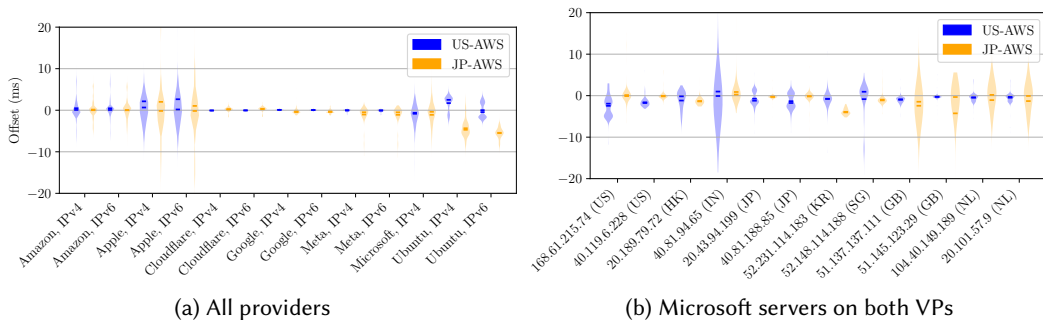


Fig. 9. Violin plots of offset distributions. Positive offset indicates that the polled server’s clock is ahead of our reference clock.

5.4 Impact of Vantage Point Location and Measurement Timing on Accuracy

Next, we investigate changes in the accuracy of the providers when measuring from other continents (North America and Asia) in June 2026, a full year after the initial measurement. We similar results as before, with the anycast providers offering very close-by servers and sub-millisecond accurate time (Appendix B). Google and Meta have nearby anycast sites in the US (<5ms RTT avg.), but not in Japan (>30ms), with Ubuntu’s time servers also being far from Japanese clients (≈ 200 ms). We see significant packet loss for Apple (>4% each measurement) and Microsoft (>2%, but not for the first measurement). Amazon does not suffer from this, despite also using DNS-based replication (§3). Lastly, some of Microsoft’s servers run at stratum 4 during our last measurement, and their accuracy improved significantly (Figure 9b).

5.5 Time services profiling

Next we determine other characteristics of the NTP services providers based on the responses we collected.

NTP version and IPv6 support: As shown in Table 6, all services deploy NTPv4 – except for Microsoft, which deploys NPTv3, standardized in 1992 [47]. Moreover, Microsoft is the only provider which does not support IPv6.

DNSSEC and RPKI support: Domain Name System Security Extensions (DNSSEC [4]) adds cryptographic signatures to DNS data to ensure its authenticity and integrity, protecting against spoofing and cache poisoning, which can protect NTP clients and others from man-in-the-middle attacks. It is used by DNS resolvers. Resource Public Key Infrastructure (RPKI [39]) provides cryptographic signatures that determine if an IP prefix can be announced by any given AS, which help to prevent BGP hijacks from false route announcements. Table 6 shows Apple and Ubuntu do not use RPKI, while DNSSEC is only used by Cloudflare.

NTS support: As of the writing of this paper, only Cloudflare and Ubuntu support NTS – Ubuntu since Oct. 2025 [43].

Stratum: The time service providers use a variety of strata (Table 6), from 1 to 4. Only Google, Meta and Apple (80% of servers) are stratum 1, although we could not measure all anycast sites of the first two. The remaining are all stratum 2 to 4 (Table 6).

Reference IDs: NTP responses including the reference ID field, describing their source of time. We classify the reference IDs obtained and show them in Figure 11.

Daemon	Version	NTP Version							Q.	
		v0	v1	v2	v3	v4	v5	v6		v7
chronyd [18]	4.7	-	v1	v2	v3	v4	-	-	-	Response
ntpd [59]	4.2.8	-	v1	v2	v3	v4	v5	v6	v7	
ntpd-rs [68]	1.6.0	-	-	-	v3	v4	-	-	-	
ntpsec [64]	1.2.4	-	v1	v2	v3	v4	-	-	-	
openntpd [9]	6.8p1	-	v1	v2	v3	v4	v5	v6	v7	

Table 7. NTP version of queries and responses.

Provider	NTP Version							Query	
	v0	v1	v2	v3	v4	v5	v6		v7
Apple	v0	v1	v2	v3	v4	v4	v4	v4	Response
	-	v1	v2	v3	v4	v5	v6	v7	
Amazon	-	v1	v2	v3	v4	-	-	-	
Cloudflare	-	v1	v2	v3	v4	-	-	-	
Google	-	v1	v2	v3	v4	v5	v6	v7	
Meta	-	v1	v2	v3	v4	-	-	-	
Microsoft	-	v3	v3	v3	v3	-	-	-	
Ubuntu	-	v1	v2	v3	v4	-	-	-	

Table 8. NTP version responses per provider

We find that 5 of the 7 providers return reference ID values that do not allow the users to figure out the upstream source. Google, for instance, returns “GOOG”, Meta returns Airport codes (LHR6), while Cloudflare, AWS, Microsoft return private IP addresses. Only Apple and Ubuntu return references that provide information about their source. However, Meta has published information online about their time services, including the time source: GNSS [65]. Apple’s stratum 1 servers all return descriptive strings, the majority returning GPSs (for: GPS via Shared Memory), and few returning SHM or MRS (for Shared Memory and likely Multi-reference Source), the two of which are generic technologies and do not describe a specific time source [33]. Apple’s stratum 2 servers return IP addresses of Apple stratum 1 servers.

Ubuntu, in turn, provides the IP addresses of their upstream sources. One of Ubuntu’s servers is synchronized with one of Apple’s stratum 1 servers, thus creating a dependency between both services. Another Ubuntu server synchronized with NIST’s server at the time. *Server software fingerprinting*: We attempt to fingerprint the software used by each provider by analyzing the NTP version field returned in the responses. First, we run the NTP server software shown in Table 7 varying the version field in the NTP request, and we observe the responses in order to obtain fingerprints of common server software.

Then, we send NTP queries to each service provide varying the NTP version field. Table 8 shows the results. Our results show that both Apple (in part) and Google may be using ntpd or openntpd. Amazon, Cloudflare, Meta and Ubuntu use software that behaves similar to Chrony and ntpsec, and there are public reports of Meta using Chrony [65]. We could not find matches for Microsoft and part of Apple’s servers.

6 Discussion and recommendations

Despite providing services for billions of devices daily, so far the large time service providers have escaped the level of scrutiny that the NTP Pool has faced. We fill the void in this paper. This is more pressing given most clients never change their NTP settings, and therefore billions of devices are served by less than a dozen of time services.

Architecture: we show a large variety in deployment architecture of time service providers – the most notable difference being the use of IP anycast (3 of the 7 providers we evaluate employ anycast). The choice depends on the operator’s preference – unicast servers combined with DNS replication give a more fine-grained control over load balancing clients.

Client mapping: previous studies have pointed out issues with the strict client/server mapping of the NTP Pool [55] – a client must be mapped to a server in the same country or continent. The non-anycast providers we evaluate do not exhibit these shortcomings: Apple, Microsoft, and AWS use more flexible methods to do such mappings. We show how Microsoft violates NTP best practices, serving half of our VPs with a single time server.

Accuracy: we show that all services provide good accuracy, with only some of Microsoft’s servers being significantly out-of-sync during part of our measurements. Considering that 50% of our VPs are served by a single server, these Microsoft clients would have no additional time servers available to notice incorrect offsets. This also demonstrates the need for longer measurements from more VPs with reference clocks for complete accuracy studies.

Recommendations: We recommend that all service providers implement NTS, DNSSEC, and RPKI to mitigate a broad range of attacks. At the time of writing, only Cloudflare supports NTS, DNSSEC, and RPKI in their time services. Alongside, these vendors must deploy robust client software that supports NTS by default, to prevent man-in-the-middle attacks.

We recommend Microsoft to adhere to the NTP best practices [70] by serving clients with multiple NTP servers; currently, roughly 50% of our vantage points that use Microsoft are served by a single time source.

Our measurements show that even large time providers can be out-of-sync. We therefore recommend that Microsoft and other non-stratum-1 providers deploy services that derive time from diverse reference sources, include non-GNSS inputs, and implement holdover mechanisms to maintain stability during outages.

7 Ethics and Disclosure

Our paper has two ethical concerns: minimizing the impact of our measurements and notifying vendors of inaccuracies and issues we have found. Most of our measurements uses RIPE Atlas as and query at low frequency (every 10min), which cannot stress the authoritative DNS servers or NTP servers that we measure – they are provisioned to answer to billions of clients.

Secondly, we found inconsistencies in some vendors, and reported them following coordinated vulnerability disclosure guidelines [57]. We have reported to Apple that one of its time servers was sending stale reference timestamps (but precise timestamps, so no impact on clock synchronization), and they have fixed the issue. We also reported to Microsoft that their DNS servers return only one NTP server most clients, which violates RFC8633 [70]. We have later informed them that we will make this information publicly 45 days from the original notification. We also notified Meta that their domain names return 1 IP address only – and it would be better if it would return multiple. We have not obtained a response. We then notified them again about the public disclosure of our findings.

8 Related Work

The closest study to ours investigates the NTP Pool architecture, client/server mapping [55], and accuracy from a subset of RIPE Atlas VPs. Different than this, we analyze seven providers and carry out accuracy experiments from two VPs synced with reference sources. [77] analyzed some of the same vendors (Apple, Canonical, Google and Microsoft next to Alibaba and NIST), but did not focus on them, leaving many questions unanswered. Our results show that the infrastructure has, in addition, changed significantly: We find more IP addresses for Apple (101 vs. 8) and Microsoft (12 vs. 9), with Apple using a geographical DNS mapping. Google’s servers now run at stratum 1 (from 2), while Microsoft’s changed to 3 (from 2). Since Microsoft’s upstreams are private IP addresses, they might still rely on NIST upstreams. The Ubuntu OS is now using Chrony instead of timesyncd [43]. Many aspects remain the same: Apple still has stratum 1 and 2 servers relying

mostly on GPS, Ubuntu has the same number of IP addresses as before, Google’s DNS configuration did not change, and Microsoft has no reverse pointers for the NTP servers. [77] describes an outage of Microsoft’s servers, while we measured subpar accuracy.

NTP Pool: The NTP Pool has been studied extensively [36, 55, 67, 77], but so far the same analysis has not been done for other large time service providers, despite them being the default providers for billions of devices. We cover them in this paper.

IP-wide scans: multiple have focussed on mapping NTP servers in the entire IP space [30, 51, 52, 77]. Considering the user base of the largest time providers, we argue that it is more important to focus on scrutinizing these – which include those in this paper and the NTP Pool, as previously shown [55].

Accuracy: A previous study has used GPS-synced VPs to measure accuracy of NTP servers [13]. However, they focus on different servers and mostly on accuracy.

NTP security: Protocol vulnerabilities have also been investigated [7, 40–42, 82], demonstrating how NTP clients can be susceptible to malicious time servers [20, 23] or to attacks by third parties [40], often focussing on broadcast-mode availability [40, 82]. Moreover, usage of NTP for amplification of distributed denial-of-service (DDoS) attacks [19] has been covered, as well as off-path attacks using DNS cache poisoning or BGP hijacking [34, 40].

9 Conclusion

NTP is a fundamental but often overlooked service on the Internet. Compared to the DNS, NTP services are significantly more centralized, as most devices use default (and sometimes hardcoded) configuration settings. We focussed on the time services of major OS and device vendors (including Windows, Ubuntu Linux, Android, macOS/iOS).

Our results show a large diversity in terms of architecture and services offered, and a case where three time servers from one provider (Microsoft) were out-of-sync. Our hope is that this study brings attention to a rather overlooked aspect of the NTP ecosystem, and can motivate vendors to improve their services by deploying servers that derive time from multiple reference sources, and that support NTS by default, as well DNSSEC and RPKI.

References

- [1] S. Alexander and R. Droms. 1997. *DHCP Options and BOOTP Vendor Extensions*. RFC 2132. IETF. <http://tools.ietf.org/rfc/rfc2132.txt>
- [2] Amazon Web Services. 2024. *Precision clock and time synchronization on your EC2 instance*. <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/set-time.html> Accessed: 2025-10-27.
- [3] Apple Inc. 2024. *Apple Reports First Quarter Results*. <https://www.apple.com/newsroom/2024/02/apple-reports-first-quarter-results/>. Accessed: 2025-02-18.
- [4] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. 2005. *DNS Security Introduction and Requirements*. RFC 4033. IETF. <http://tools.ietf.org/rfc/rfc4033.txt>
- [5] bgptools. 2025. *anycast-prefixes*. <https://github.com/bgptools/anycast-prefixes> Accessed: 2025-02-20.
- [6] Leo Bicknell. 2012. *NTP Issues Today*. <https://mailman.nanog.org/pipermail/nanog/2012-November/053449.html>.
- [7] M. Bishop. 1990. A security analysis of the NTP protocol version 2. In *[1990] Proceedings of the Sixth Annual Computer Security Applications Conference*. 20–29. <https://doi.org/10.1109/CSAC.1990.143746>
- [8] Russel Brandon. 2021. *Google announces that Android has over 3 billion active devices*. <https://www.theverge.com/2021/5/18/22440813/android-devices-active-number-smartphones-google-2021> Accessed: 2025-02-18.
- [9] Henning Brauer. 2022. *OpenNTPd*. <https://www.openntpd.org> Accessed: 2025-05-13.
- [10] Physikalisch-Technische Bundesanstalt. 2025. *Questions about Time*. <https://www.ptb.de/cms/en/ptb/fachabteilungen/abt4/fb-44/fragenzurzeit/fragenzurzeit07.html> Accessed: 2025-05-01.
- [11] Physikalisch-Technische Bundesanstalt. n.d. *DCF77*. <https://www.ptb.de/cms/en/ptb/fachabteilungen/abt4/fb-44/ag-442/dissemination-of-legal-time/DCF77.html>. Accessed: 2025-11-06.
- [12] R. Bush and R. Austein. 2013. *The Resource Public Key Infrastructure (RPKI) to Router Protocol*. RFC 6810. IETF. <http://tools.ietf.org/rfc/rfc6810.txt>


- [13] Yi Cao and Darryl Veitch. 2020. Toward Trusted Time: Remote Server Vetting and the Misfiring Heart of Internet Timing. *IEEE/ACM Transactions on Networking* 28, 2 (2020), 944–956. <https://doi.org/10.1109/TNET.2020.2977024>
- [14] RIPE Network Coordination Centre. 2025. *K-Root*. <https://www.ripe.net/analyse/dns/k-root/> Accessed: 2025-03-13.
- [15] RIPE Network Coordination Centre. 2025. *RIPE Atlas Probe Archive*. <https://ftp.ripe.net/ripe/atlas/probes/archive/> Accessed: 2025-02-12.
- [16] Cloudflare. 2024. Cloudflare Time Service. <https://www.cloudflare.com/time/>.
- [17] Internet Systems Consortium. 2025. *F-Root*. <https://www.isc.org/f-root/> Accessed: 2025-03-13.
- [18] Richard Curnow and Miroslav Lichvar. 2024. *chrony - Network Time Protocol (NTP) Implementation*. <https://chrony-project.org/> Accessed: 2025-03-04.
- [19] Jakub Czyz, Michael Kallitsis, Manaf Gharaibeh, Christos Papadopoulos, Michael Bailey, and Manish Karir. 2014. Taming the 800 Pound Gorilla: The Rise and Decline of NTP DDoS Attacks. In *Proceedings of the 2014 ACM Conference on Internet Measurement Conference (Vancouver, BC, Canada) (IMC)*. ACM, 435–448. <https://doi.org/10.1145/2663716.2663717>
- [20] Omer Deutsch, Neta Rozen Schiff, Danny Dolev, and Michael Schapira. 2018. Preventing (Network) Time Travel with Chronos. In *NDSS*.
- [21] Linux PTP development team. 2025. *LinuxPTP – Precision Time Protocol (PTP) implementation for Linux*. <https://linuxptp.sourceforge.net/> Accessed: 2025-11-05.
- [22] D. Franke, D. Sibold, K. Teichel, M. Dansarie, and R. Sundblad. 2020. *Network Time Security for the Network Time Protocol*. RFC 8915. IETF. <http://tools.ietf.org/rfc/rfc8915.txt>
- [23] Tweede Golf. 2024. Manipulating time through (S)NTP. <https://tweedegolf.nl/en/blog/142/manipulating-time-through-sntp> Accessed: 2025-11-18.
- [24] Google. 2025. Network edge locations. <https://cloud.google.com/about/locations/>. [Accessed 15-05-2025].
- [25] Bjørn Hansen. 2023. Minor New Features on the website. <https://community.ntppool.org/t/minor-new-features-on-the-website/2947/8>.
- [26] Remi Hendriks, Matthew Luckie, Mattijs Jonker, Raffaele Sommese, and Roland van Rijswijk-Deij. 2025. LACeS: An Open, Fast, Responsible and Efficient Longitudinal Anycast Census System. In *Proceedings of the 2025 ACM Internet Measurement Conference (IMC '25)*. 445–461. <https://doi.org/10.1145/3730567.3764484>
- [27] Remi Hendriks, Matthew Luckie, Mattijs Jonker, Raffaele Sommese, and Roland van Rijswijk-Deij. 2025. MAnycast Reloaded: a Tool for an Open, Fast, Responsible and Efficient Daily Anycast Census. arXiv:2503.20554 [cs.NI] <https://arxiv.org/abs/2503.20554>
- [28] Philip Homburg. 2015. NTP Measurements with RIPE Atlas. https://labs.ripe.net/author/philip_homburg/ntp-measurements-with-ripe-atlas/.
- [29] Nate Hopper. 2022. The Thorny Problem of Keeping the Internet’s Time. *The New Yorker* (Sept. 30 2022). <https://www.newyorker.com/tech/annals-of-technology/the-thorny-problem-of-keeping-the-internets-time>
- [30] Zhentian Huang, Shuai Wang, Li Chen, Dan Li, and Yilun Liu. 2025. Measuring the Time Source Vulnerabilities in the NTP Ecosystem. In *Proceedings of the 2025 ACM Internet Measurement Conference (USA) (IMC '25)*. Association for Computing Machinery, New York, NY, USA, 16 pages.
- [31] IEEE. 2020. IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems. *IEEE Std 1588-2019 (Revision of IEEE Std 1588-2008)* (2020), 1–499. <https://doi.org/10.1109/IEEESTD.2020.9120376>
- [32] IPinfo.io. 2025. *IPinfo - Trusted IP Data Provider*. <https://ipinfo.io> Accessed: 2025-02-12.
- [33] Andreja Jarc. 2017. *Multi-Reference Clock*. <https://blog.meinbergglobal.com/2017/11/16/multi-reference-clock/> Accessed: 2025-11-10.
- [34] Philipp Jeitner, Haya Shulman, and Michael Waidner. 2020. The Impact of DNS Insecurity on Time. In *2020 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 266–277. <https://doi.org/10.1109/DSN48063.2020.00043>
- [35] S. Konjerla, G. C. M. Moura, G. Smaragdakis, and T. Dittrich. 2025. *Are NTP Clients Always Right? Evaluating NTP Clients under Normal and Attack Scenarios*. Technical Report SIDN Labs Technical Report 2025-10-16. SIDN Labs and TU Delft, Arnhem, The Netherlands. https://www.sidnlabs.nl/downloads/7M1bz1otGu0D7hqr0hRT2c/3bc09536c3bc87f63b80d66944abc1d9/ntp-clients-tech_report-20251016.pdf Updated October 30, 2025.
- [36] Jonghoon Kwon, Jeonggyu Song, Junbeom Hur, and Adrian Perrig. 2023. Did the Shark Eat the Watchdog in the NTP Pool? Deceiving the NTP Pool’s Monitoring System. In *30th USENIX Security Symposium*. <https://www.usenix.org/conference/usenixsecurity23/presentation/kwon>
- [37] National Physical Laboratory. 2025. *Internet Time Service*. <https://www.npl.co.uk/products-services/time-frequency/internet-time> Accessed: 2025-05-01.
- [38] leeand00. 2019. *Is it true that iOS devices use a hard coded time.apple.com DNS for an NTP server?* <https://apple.stackexchange.com/questions/364765/is-it-true-that-ios-devices-use-a-hard-coded-time-apple-com-dns-for-an-ntp-serve> “Asked 6 years, 3 months ago” on Ask Different (Stack Exchange).

- [39] M. Lepinski and S. Kent. 2012. *An Infrastructure to Support Secure Internet Routing*. RFC 6480. IETF. <http://tools.ietf.org/rfc/rfc6480.txt>
- [40] Aanchal Malhotra, Isaac E Cohen, Erik Brakke, and Sharon Goldberg. 2016. Attacking the Network Time Protocol. In *Proceedings of the 23rd Network and Distributed System Security Symposium (NDSS 2016)* (San Diego, California).
- [41] Aanchal Malhotra and Sharon Goldberg. 2016. Attacking NTP's Authenticated Broadcast Mode. *SIGCOMM Comput. Commun. Rev.* 46, 2 (may 2016), 12–17.
- [42] Aanchal Malhotra, Matthew Van Gundy, Mayank Varia, Haydn Kennedy, Jonathan Gardner, and Sharon Goldberg. 2017. The Security of NTP's Datagram Protocol. In *Financial Cryptography and Data Security: 21st International Conference, FC 2017, Sliema, Malta, April 3-7, 2017, Revised Selected Papers 21*. Springer, 405–423.
- [43] Lukas Mårdian. 2025. PSA: Installing chrony by default, to enable Network Time Security (NTS). Ubuntu-devel mailing list. <https://lists.ubuntu.com/archives/ubuntu-devel/2025-May/043355.html> Message posted on May 22, 2025.
- [44] Mark Morowczynski. 2012. Did Your Active Directory Domain Time Just Jump To The Year 2000? <https://techcommunity.microsoft.com/t5/core-infrastructure-and-security/did-your-active-directory-domain-time-just-jump-to-the-year-2000/ba-p/255873>.
- [45] Microsoft. 2022. Windows Devices by the Numbers. <https://web.archive.org/web/20220419050729/https://news.microsoft.com/bythenumbers/en/windowsdevices>. Accessed: 2025-02-18.
- [46] Microsoft. 2025. *Traffic Manager Geographic Traffic Routing Method*. <https://learn.microsoft.com/en-us/azure/traffic-manager/traffic-manager-faqs#traffic-manager-geographic-traffic-routing-method> Accessed: 2025-05-02.
- [47] D. Mills. 1992. *Network Time Protocol (Version 3) Specification, Implementation and Analysis*. RFC 1305. IETF. <http://tools.ietf.org/rfc/rfc1305.txt>
- [48] D. Mills. 1996. *Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI*. RFC 2030. IETF. <http://tools.ietf.org/rfc/rfc2030.txt>
- [49] D. Mills, J. Martin, J. Burbank, and W. Kasch. 2010. *Network Time Protocol Version 4: Protocol and Algorithms Specification*. RFC 5905. IETF. <http://tools.ietf.org/rfc/rfc5905.txt>
- [50] David L Mills. 1991. Internet time synchronization: the network time protocol. *IEEE Transactions on Communications* 39, 10 (1991), 1482–1493. <https://doi.org/10.1109/26.103043>
- [51] David L. Mills, Ajit Thyagarjan, and Brian C. Huffman. 1997. Internet Timekeeping Around the Globe. <https://www.eecis.udel.edu/~mills/database/papers/survey5.pdf>
- [52] Nelson Minar. 1999. A Survey of the NTP Network. <https://www.eecis.udel.edu/~mills/database/reports/ntp-survey99-minar.pdf>
- [53] Faten Mkacher and Andrzej Duda. 2019. Calibrating NTP. In *2019 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS)*. 1–6. <https://doi.org/10.1109/ISPCS.2019.8886646>
- [54] Giovane C. M. Moura, Sebastian Castro, Wes Hardaker, Maarten Wullink, and Cristian Hesselman. 2020. Clouding up the Internet: How Centralized is DNS Traffic Becoming?. In *Proceedings of the ACM Internet Measurement Conference (Virtual Event, USA) (IMC '20)*. Association for Computing Machinery, New York, NY, USA, 42–49.
- [55] Giovane C. M. Moura, Marco Davids, Caspar Schutjser, Cristian Hesselman, John Heidemann, and Georgios Smaragdakis. 2024. Deep Dive into NTP Pool's Popularity and Mapping. *Proceedings of the ACM on Measurement and Analysis of Computing Systems (SIGMETRICS 2024)* 8, 1, Article 15 (feb 2024), 30 pages. <https://doi.org/10.1145/3639041>
- [56] Giovane C. M. Moura, Ricardo de O. Schmidt, John Heidemann, Wouter B. de Vries, Moritz Müller, Lan Wei, and Christian Hesselman. 2016. Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event. In *Proceedings of the ACM Internet Measurement Conference*. <https://doi.org/10.1145/2987443.2987446>
- [57] Giovane C. M. Moura and John Heidemann. 2023. Vulnerability Disclosure Considered Stressful. *SIGCOMM Comput. Commun. Rev.* 53, 2 (July 2023), 2–10. <https://doi.org/10.1145/3610381.3610383>
- [58] Giovane C. M. Moura, John Heidemann, Ricardo de O. Schmidt, and Wes Hardaker. 2019. Cache Me If You Can: Effects of DNS Time-to-Live. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Amsterdam, the Netherlands, 101–115. <https://doi.org/10.1145/3355369.3355568>
- [59] Network Time Foundation. 2025. *NTPD 4.2.8p-series*. <https://www.ntp.org/documentation/4.2.8-series/#building-and-installing-ntp> Accessed: 2025-05-13.
- [60] NIST. 2025. NIST Internet Time Service (ITS). (May. 1 2025). <https://www.nist.gov/pml/time-and-frequency-division/time-distribution/internet-time-service-its>
- [61] Yevheniya Nosyk, Qasim Lone, Yury Zhauniarovich, Carlos H. Gañán, Emile Aben, Giovane C. M. Moura, Samaneh Tajalizadehkhoo, Andrzej Duda, and Maciej Korczyński. 2023. Intercept and Inject: DNS Response Manipulation in the Wild. In *Passive and Active Measurement*, Anna Brunstrom, Marcel Flores, and Marco Fiore (Eds.). Springer Nature Switzerland, Cham, 461–478.
- [62] Andrew N. Novick and Michael A. Lombardi. 2015. Practical limitations of NTP time transfer. In *2015 Joint Conference of the IEEE International Frequency Control Symposium & the European Frequency and Time Forum*. 570–574. <https://doi.org/10.1109/FCS.2015.7138909>


- [63] NTP Pool. 2024. pool.ntp.org: the Internet cluster of NTP servers. <https://www.ntppool.org/en/>.
- [64] NTPsec project. 2022. NTPSec. <https://www.ntpsec.org> Accessed: 2025-05-13.
- [65] Oleg Obleukhov. 2020. Building a more accurate time service at Facebook scale. <https://engineering.fb.com/2020/03/18/production-engineering/ntp-service/>.
- [66] C. Partridge, T. Mendez, and W. Milliken. 1993. *Host Anycasting Service*. RFC 1546. IETF. <http://tools.ietf.org/rfc/rfc1546.txt>
- [67] Yarin Perry, Neta Rozen-Schiff, and Michael Schapira. 2021. A Devil of a Time: How Vulnerable is NTP to Malicious Timeservers?. In *Proceedings of the 28th Network and Distributed System Security Symposium (NDSS 2021)* (Virtual Conference).
- [68] Pendulum Project. [n. d.]. *ntpd-rs: A full-featured implementation of the Network Time Protocol, including NTS support*. <https://github.com/pendulum-project/ntpd-rs> GitHub repository.
- [69] The Linux Kernel Documentation Project. 2024. PTP hardware clock infrastructure for Linux. <https://docs.kernel.org/driver-api/ptp.html>. Accessed: 2025-11-05.
- [70] D. Reilly, H. Stenn, and D. Sibold. 2019. *Network Time Protocol Best Current Practices*. RFC 8633. IETF. <http://tools.ietf.org/rfc/rfc8633.txt>
- [71] E. Rescorla. 2018. *The Transport Layer Security (TLS) Protocol Version 1.3*. RFC 8446. IETF. <http://tools.ietf.org/rfc/rfc8446.txt>
- [72] RIPE NCC. 2025. RIPE Atlas Measurements. <https://atlas.ripe.net/measurements/ID>. where ID is the measurement ID: Google-android: 100787434, Meta: 100787995, Meta-v6: 100930063, Ubuntu: 100788409, Apple-ZA-map: 87282270, Microsoft-map: 86743242, Apple-map1: 86743255, Apple-map2: 86743251, Ubuntu-map: 87076806, AWS-map: 87077187, AWS-map-v6: 90364265, ms-server1-detection: 87119082, ubuntu-local-server: 90982800, apple-wrong-refts: 89211992, apple-local-servers: [89183871-89183874,89183992].
- [73] RIPE NCC. 2026. RIPE Atlas Measurement IDs. <https://atlas.ripe.net/measurements/ID>. , where ID is the experiment ID: Apple-long: 160722842, Microsoft-long: 160723781, AWS-long: 160723821.
- [74] RIPE NCC Staff. 2015. RIPE Atlas: A Global Internet Measurement Network. *Internet Protocol Journal (IPJ)* 18, 3 (Sep 2015), 2–26.
- [75] RIPE Network Coordination Centre. 2025. RIPE Atlas. <https://atlas.ripe.net>.
- [76] Root Server Operators. 2025. Root DNS. <http://root-servers.org/>.
- [77] Teemu Ryttilahti, Dennis Tatang, Janosch Köpper, and Thorsten Holz. 2018. Masters of Time: An Overview of the NTP Ecosystem. In *2018 IEEE European Symposium on Security and Privacy (EuroS P)*. 122–136. <https://doi.org/10.1109/EuroSP.2018.00017>
- [78] Amazon Web Services. 2022. *Amazon Time Sync is now available over the internet as a public NTP service*. <https://aws.amazon.com/about-aws/whats-new/2022/11/amazon-time-sync-internet-public-ntp-service/> Accessed: 2025-03-11.
- [79] SIDN Labs. 2025. TimeNL. https://time.nl/index_en.html.
- [80] Raffaele Sommese, Leandro Bertholdo, Gautam Akiwate, Mattijs Jonker, Roland van Rijswijk-Deij, Alberto Dainotti, KC Claffy, and Anna Sperotto. 2025. *LACes Anycast Census*. <https://manycast.net/> Accessed: 2025-03-06.
- [81] Raffaele Sommese and Remi Hendriks. 2021. LACes Anycast Census Data Repository. <https://github.com/ut-dacs/Anycast-Census>.
- [82] Nikhil Tripathi and Neminath Hubballi. 2021. Preventing time synchronization in NTP broadcast mode. *Computers & Security* 102 (2021), 102135. <https://doi.org/10.1016/j.cose.2020.102135>
- [83] Kevin Vermeulen, Ege Gurmericliler, Italo Cunha, David Choffnes, and Ethan Katz-Bassett. 2022. Internet Scale Reverse Traceroute. In *Proceedings of the 22nd ACM Internet Measurement Conference (Nice, France) (IMC '22)*. Association for Computing Machinery, New York, NY, USA, 694–715. <https://doi.org/10.1145/3517745.3561422>
- [84] Lan Wei and John Heidemann. 2017. Does Anycast Hang Up On You?. In *IEEE Network Traffic Monitoring and Analysis Conference*. IEEE, Dublin, Ireland, 9. <https://doi.org/10.23919/TMA.2017.8002905>

A Extra Graphs

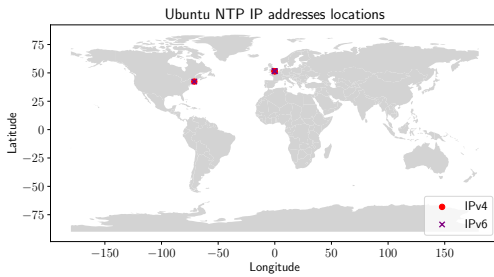
Figure 10a shows the results from §4.3 for time-osx.g.applimg.com, Apple’s second domain name. Figure 10b shows that all of Ubuntu’s 4 IPv4 addresses are included in every DNS response without any load balancing or mapping. Figure 10c shows the geolocation of Ubuntu’s IP addresses: they are in only 2 different locations. Figure 10d shows the mappings of RIPE Atlas probes’ countries to Apple NTP servers for time-osx.g.aapplimg.com. It follows the same mapping (and the same servers) as Figure 4b. Figure 10e shows the measured average absolute time offsets per provider.




(a) Number of responses for each NTP server for time-osx.g.aapling.com



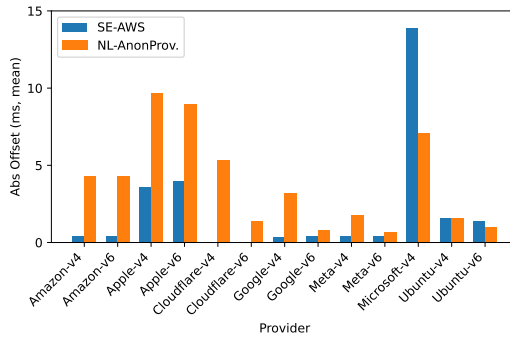
(b) Number of responses for each NTP server for ntp.ubuntu.com



(c) Geographical location of Ubuntu NTP server IP addresses



(d) Sankey diagram of country of Atlas Probes VPs (left side) and Apple's IPv4 NTP servers' countries (right side) for time-osx.g.aapling.com. Dataset: Apple-map-2[72]




(e) Average absolute offset from the NTP-measurement


Figure 10 shows the number of responses for Apple's US and EU-based NTP servers. Figure 11 shows the reference IDs reported by each provider's time servers.

B Extra Tables

Table 14 gives numerical statistics on the out-of-sync servers we found; cf. Figure 8. Table 9 gives a detailed overview of the DNS responses for the Microsoft, Apple, and AWS time servers. Table 10 shows the accuracy experiment results for NL-SIDN.



(a) Apple (time.g.aaplimg.com – US Servers)



(b) Apple (time.g.aaplimg.com – EU Servers)

Fig. 10. Number of responses for each NTP server for Apple US and EU Servers




Fig. 11. Sankey plot of references (time sources) used by each provider’s servers (SE-AWS).

Query name	Microsoft twc.trafficmanager.net A	Apple1 time.g.aaplimg.com A	Apple2 time-osx.g.aaplimg.com A	AWS time.aws.com
Dataset	Microsoft-map	Apple-map-1	Apple-map-2	AWS-map
Auth server	13.107.222.240	17.253.206.8	17.253.201.8	205.251.193.84
Total Probes	9,231	9,205	9,211	9,068
without valid response	221	211	223	129
with valid responses	9,010	8,994	8,988	8,939
only with valid responses	8,270	8,066	8,248	8,284
with valid and invalid responses	740	928	740	655
Total countries	171	171	172	172
without valid response	52	52	53	41
with valid responses	170	170	171	171
only with valid responses	167	163	164	165
with valid and invalid responses	170	170	171	171
Total ASes	3,143	3,142	3,143	3,115
without valid response	158	152	158	104
with valid responses	3,094	3,094	3,093	3,081
only with valid responses	2,865	2,874	2,921	2,905
with valid and invalid responses	487	489	394	424
DNS queries	663,637	661,834	662,165	652,105
valid response (RCODE 0)	642,617	642,042	641,436	640,253
Timeout/ network error	11,945	12,896	12,080	11,754
SERVFAIL (RCODE 2)	4,378	2,670	4337	20
REFUSED (RCODE 5)	4,697	4,210	4,312	78
malformed response	0	16	0	0
	Time Servers			
Unique NTP servers	15	57	55	93
Valid	12	53	53	90
Invalid	3	4	2	3
ASes NTP Servers	1 (AS8075)	2 (AS6185 and AS714)	2 (AS6185 and AS714)	2 (AS14618 and AS16509)

Table 9. Mapping Time Service Networks. Datasets: [72]

Provider	IPv	#IP Addr.	RPKI	DNS-SEC	#Req.	#Resp.	No Resp. #	No Resp. %	Offset mean	Abs. Offset 90th %ile	RTT mean	NTP version	NTS	Stratum
Amazon	v4	88	Yes	No	25 344	25 329	15	0.06%	2.55 ms	10.16 ms	151.36 ms	v4	No	4
	v6	88	Yes	No	25 344	25 327	17	0.07%	2.59 ms	9.92 ms	151.02 ms	v4	No	4
Apple	v4	46	No	No	13 248	12 366	882	6.66%	3.71 ms	20.90 ms	174.68 ms	v4	No	1/2 (77%/23%)
	v6	45	No	No	12 960	11 630	1 330	10.26%	-2.09 ms	20.33 ms	155.06 ms	v4	No	1/2 (79%/21%)
Cloudflare	v4	2	Yes	Yes	576	576	0	0.00%	-4.82 ms	4.31 ms	15.18 ms	v4	Yes	3
	v6	2	Yes	Yes	576	576	0	0.00%	-0.68 ms	0.81 ms	6.73 ms	v4	Yes	3
Google	v4	4	Yes	No	1 152	1 152	0	0.00%	-3.05 ms	1.33 ms	14.37 ms	v4	No	1
	v6	4	Yes	No	1 152	1 152	0	0.00%	-0.62 ms	0.30 ms	9.39 ms	v4	No	1
Meta	v4	5	Yes	No	1 440	1 311	129	8.96%	-1.65 ms	3.92 ms	16.43 ms	v4	No	1
	v6	5	Yes	No	1 440	1 440	0	0.00%	-0.46 ms	1.38 ms	14.83 ms	v4	No	1
Microsoft	v4	12	Yes	No	3 456	3 303	153	4.43%	-6.37 ms	5.39 ms	130.98 ms	v3	No	3
Ubuntu	v4	4	No	No	1 152	1 146	6	0.52%	-1.24 ms	3.52 ms	26.13 ms	v4	Yes	2
	v6	3	No	No	864	863	1	0.12%	0.52 ms	0.98 ms	9.73 ms	v4	Yes	2

Table 10. Statistics from the measurement results of NL-SIDN. Anycasted services are highlighted in orange.

Provider	IPv	#IP Addr.	RPKI	DNS-SEC	#Req.	#Resp.	No Resp. #	No Resp. %	Offset mean	Abs. Offset 90th %ile	RTT mean	NTP version	NTS	Stratum
Amazon	v4	75	Yes	No	21 599	21 585	14	0.06%	0.48 ms	4.86 ms	128.65 ms	v4	No	4
	v6	75	Yes	No	21 600	21 592	8	0.04%	0.46 ms	4.93 ms	128.66 ms	v4	No	4
Apple	v4	27	No	No	7 776	7 404	372	4.78%	2.14 ms	4.54 ms	141.00 ms	v4	No	1/2 (72%/28%)
	v6	27	No	No	7 776	7 151	625	8.04%	2.64 ms	5.41 ms	123.54 ms	v4	No	1/2 (75%/25%)
Cloudflare	v4	2	Yes	Yes	576	576	0	0.00%	-0.07 ms	0.18 ms	1.42 ms	v4	Yes	3
	v6	2	Yes	Yes	576	576	0	0.00%	-0.05 ms	0.17 ms	1.38 ms	v4	Yes	3
Google	v4	4	Yes	No	1 152	1 152	0	0.00%	0.10 ms	0.23 ms	2.99 ms	v4	No	1
	v6	4	Yes	No	1 152	1 152	0	0.00%	0.05 ms	0.21 ms	2.91 ms	v4	No	1
Meta	v4	5	Yes	No	1 440	1 440	0	0.00%	-0.07 ms	0.47 ms	4.97 ms	v4	No	1
	v6	5	Yes	No	1 440	1 440	0	0.00%	-0.10 ms	0.52 ms	4.94 ms	v4	No	1
Microsoft	v4	12	Yes	No	3 456	3 372	84	2.43%	-0.85 ms	3.52 ms	126.76 ms	v3	No	4/3 (58%/42%)
Ubuntu	v4	4	No	No	1 152	1 102	50	4.34%	1.71 ms	3.17 ms	66.28 ms	v4	Yes	2
	v6	3	No	No	864	864	0	0.00%	0.02 ms	2.58 ms	85.15 ms	v4	Yes	2

Table 11. Statistics from the measurement results of our US-AWS. Anycasted services are highlighted in orange.

Provider	IPv	#IP Addr.	RPKI	DNS-SEC	#Req.	#Resp.	No Resp. #	No Resp. %	Offset mean	Abs. Offset 90th %ile	RTT mean	NTP version	NTS	Stratum
Amazon	v4	75	Yes	No	21 600	21 581	19	0.09%	0.16 ms	4.17 ms	157.65 ms	v4	No	4
	v6	75	Yes	No	21 600	21 583	17	0.08%	0.00 ms	3.92 ms	157.94 ms	v4	No	4
Apple	v4	27	No	No	7 776	7 425	351	4.51%	2.02 ms	11.02 ms	147.99 ms	v4	No	1/2 (72%/28%)
	v6	27	No	No	7 776	7 169	607	7.81%	1.04 ms	13.37 ms	136.43 ms	v4	No	1/2 (75%/25%)
Cloudflare	v4	2	Yes	Yes	576	576	0	0.00%	0.20 ms	1.03 ms	1.92 ms	v4	Yes	3
	v6	2	Yes	Yes	576	576	0	0.00%	0.27 ms	1.03 ms	1.72 ms	v4	Yes	3
Google	v4	4	Yes	No	1 152	1 152	0	0.00%	-0.32 ms	0.98 ms	35.28 ms	v4	No	1
	v6	4	Yes	No	1 152	1 152	0	0.00%	-0.31 ms	0.95 ms	35.01 ms	v4	No	1
Meta	v4	5	Yes	No	1 440	1 440	0	0.00%	-1.03 ms	2.15 ms	43.51 ms	v4	No	1
	v6	5	Yes	No	1 440	1 440	0	0.00%	-1.10 ms	2.24 ms	43.49 ms	v4	No	1
Microsoft	v4	12	Yes	No	3 456	3 312	144	4.17%	-1.11 ms	4.74 ms	127.26 ms	v3	No	4/3 (59%/41%)
Ubuntu	v4	4	No	No	1 152	1 142	10	0.87%	-4.33 ms	6.67 ms	204.01 ms	v4	Yes	2
	v6	3	No	No	864	864	0	0.00%	-5.52 ms	7.37 ms	215.30 ms	v4	Yes	2

Table 12. Statistics from the measurement results of our JP-AWS. Anycasted services are highlighted in orange.

C Offset Bounds and the Effect of Path Asymmetry

Figure 1 illustrates the standard NTP four-timestamp exchange, where T_1 and T_4 are recorded by the client, and T_2 and T_3 are recorded by the server. If θ denotes the true offset of the server clock relative to the client, and δ is the *measured* RTT, i.e., $\delta = (T_4 - T_1) - (T_3 - T_2)$.

Symmetric paths. If communication delays are known and constant (d_1, d_2), the offset can be calculated from either one-way delay:

$$\theta = T_2 - T_1 - d_1 = -(T_4 - T_3 - d_2)$$

	Apple	Microsoft	AWS
Probes	8,929	8,930	8,930
Queries	2,835,666	2,793,836	2,836,777
Responses	2,714,279	2,675,008	2,714,153
NTP servers	52	12	86
Dataset ID	160722842	160723781	160723821

Table 13. Measurement summary across the three target services. Datasets: [73]

Server	104.40.149.189	20.101.57.9	51.137.137.111
# Mismatches	20	137	42
Largest Mismatch	-4.20ms	-39.08ms	-9.16ms
Mean RTT	22.31ms	22.93ms	31.57ms
Abs. Offset Q1	11.35ms	22.92ms	16.38ms
Abs. Offset Q4	13.14ms	33.74ms	18.98ms
Largest Offset	-15.61ms	-50.97ms	-24.51ms

Table 14. Statistics about out-of-sync servers (SE-AWS): Mismatch = $T_2 - T_1$ if $T_1 > T_2$ else $T_3 - T_4$ if $T_3 > T_4$ else 0, Q1 and Q4: first and last quartile

If d_1 and d_2 are unknown, but assumed equal, then we can calculate them using the round-trip time $d_1 = d_2 = \frac{\delta}{2} = \frac{d_1+d_2}{2} = \frac{(T_4-T_1)-(T_3-T_2)}{2} = \frac{(T_4-T_3)+(T_2-T_1)}{2}$, yielding

$$\theta_{\text{NTP}} = T_2 - T_1 - d_1 = -(T_4 - T_3 - d_2) = T_2 - T_1 - \frac{(T_4 - T_3) + (T_2 - T_1)}{2} = \frac{(T_2 - T_1) - (T_4 - T_3)}{2}$$

which is the classical NTP offset calculation, giving correct/exact offsets for equal delays.

Asymmetric paths. In practice, delays are neither equal nor constant, but depend on Internet routing. Asymmetry is common, i.e. $d_1 \neq d_2$, introducing a bias in the measured offset, i.e., a difference between θ and θ_{NTP} :

$$\begin{aligned} \theta &= \frac{(\theta + \theta)}{2} = \frac{T_2 - T_1 - d_1 - (T_4 - T_3 - d_2)}{2} = \frac{(T_2 - T_1) - (T_4 - T_3) + d_2 - d_1}{2} \\ &= \frac{(T_2 - T_1) - (T_4 - T_3) + d_2 - d_1}{2} = \frac{(T_2 - T_1) - (T_4 - T_3)}{2} + \frac{d_2 - d_1}{2} = \theta_{\text{NTP}} + \frac{d_2 - d_1}{2} \end{aligned}$$

Because $d_1, d_2 \geq 0$, it follows that $-(d_2 - d_1) \leq d_2 - d_1 \leq d_2 + d_1$, with those bounds occurring when one delay is 0. Therefore the corresponding bounds expand to $\frac{d_1+d_2}{2} \leq \theta - \theta_{\text{NTP}} \leq \frac{d_1+d_2}{2}$, i.e. $|\theta - \theta_{\text{NTP}}| \leq \frac{\delta}{2}$.

D Extra Measurements for Mapping Time Services

To verify that our 8-hour campaign observes the full set of NTP servers behind each provider, we ran a longer measurement using RIPE Atlas spanning over 60 days, covering April and May 2026. We restrict this campaign to the providers that do not rely on IP anycast (Apple, AWS, and Microsoft), as anycast already obscures the mapping between probes and individual servers. The datasets are available at [73].

Table 13 summarizes the results. Despite probing over a window roughly 180× longer than the one in §3.2, we recover the similar server counts—52 for Apple, 12 for Microsoft, and 86 for AWS—confirming that 8 hours already suffice to enumerate the complete set.