Internet Trends as Seen from IIJ Infrastructure —2024

Internet services provider IIJ operates some of the largest network and server infrastructure in Japan. Each year, we analyze a year's worth of Internet trends based on information obtained through the operation of this infrastructure and report on it in the IIR. This year, we again analyze changes in trends over the past year from the perspective of BGP routes, DNS query analysis, IPv6, and mobile.

Topic 1

BGP and Routes

We start by looking at IPv4 full-route information advertised by our network to other organizations (Table 1) and the number of unique IPv4 addresses contained in the IPv4 full-route information (Table 3).

The total number of routes reached just under 950,000. While the increase was more than double last year's, it was still the second smallest increase in the past decade (see Figure 1), and that downtrend appears to be ongoing. This year, we observed decreases in all routes with /10–/20 prefixes. This is likely due to the subdivision of address blocks for transfer purposes still being quite prominent. The number of unique IPv4 addresses fell by over 22 million (0.6%), an even larger decline than last year. Combined with last year's decrease, this represents a loss equivalent to about two /8 blocks.

Next, we look at IPv6 full-route information (Table 2) and the number of unique IPv6 /64 blocks in the IPv6 full-route information (Table 3).

Table 1: Number of Routes by Prefix Length for Full IPv4 Routes

Date	/8	/9	/10	/11	/12	/13	/14	/15	/16	/17	/18	/19	/20	/21	/22	/23	/24	total
Sep. 2015	18	13	36	96	261	500	999	1731	12863	7190	12317	25485	35904	38572	60900	52904	301381	551170
Sep. 2016	16	13	36	101	267	515	1050	1767	13106	7782	12917	25229	38459	40066	67270	58965	335884	603443
Sep. 2017	15	13	36	104	284	552	1047	1861	13391	7619	13385	24672	38704	41630	78779	64549	367474	654115
Sep. 2018	14	11	36	99	292	567	1094	1891	13325	7906	13771	25307	39408	45578	88476	72030	400488	710293
Sep. 2019	10	11	37	98	288	573	1142	1914	13243	7999	13730	25531	40128	47248	95983	77581	438926	764442
Sep. 2020	9	11	39	100	286	576	1172	1932	13438	8251	14003	25800	40821	49108	101799	84773	473899	816017
Sep. 2021	16	13	41	101	303	589	1191	2007	13408	8231	13934	25276	41915	50664	106763	91436	497703	853591
Sep. 2022	16	13	39	101	298	592	1208	2064	13502	8292	13909	25051	43972	52203	109071	96909	536520	903760
Sep. 2023	16	14	39	102	298	577	1196	2064	13490	8245	13809	25059	43863	51012	109514	98178	550621	918097
Sep. 2024	16	16	37	93	295	573	1165	2059	13224	8220	13718	24624	43786	51827	111483	99239	579274	949649

Table 2: Number of Routes by Prefix Length for Full IPv6 Routes

Date	/16-/28	/29	/30-/31	/32	/33-/39	/40	/41-/43	/44	/45-/47	/48	total
Sep. 2015	142	771	168	6846	1808	1150	386	990	648	10570	23479
Sep. 2016	153	1294	216	8110	3092	1445	371	1492	1006	14291	31470
Sep. 2017	158	1757	256	9089	3588	2117	580	1999	1983	18347	39874
Sep. 2018	168	2279	328	10897	4828	2940	906	4015	2270	24616	53247
Sep. 2019	192	2671	606	12664	6914	3870	1566	4590	4165	34224	71462
Sep. 2020	205	3164	641	14520	9063	4815	2663	5501	4562	45160	90294
Sep. 2021	223	3628	705	20650	13050	10233	4170	11545	5204	61024	130432
Sep. 2022	298	4247	895	21926	15147	12509	4108	13840	6994	73244	153208
Sep. 2023	316	4357	923	23228	17427	14828	5518	16453	9579	86881	179510
Sep. 2024	322	5360	934	24739	20198	17657	4672	19418	12470	95628	201398

The total number of routes surpassed 200,000. While the increase was smaller than last year, it remained at about the same level as two years ago. Excluding those in the "/41–/43" column, route numbers have been on the rise since we started these periodic observations, but the first ever four-digit increase in /29 routes is noteworthy here. The number of unique /64 blocks also increased by nearly 30%, indicating that the IPv6 rollout and expansion of IPv6 networks continues to progress. As a side note unrelated to BGP, a larger IPv6 address block (/20) has been added for documentation purposes (RFC 9637), as the original block (/32) was deemed inadequate for examples of modern large-scale networks.

Lastly, let's also look at IPv4/IPv6 full-route Origin AS figures (Table 4). In the past year, ARIN and LACNIC were each allocated an additional 1024 32-bit only ASNs.

Eight years have passed since the IANA's 16-bit ASN pool was depleted in 2016. While the nine-year-straight drop in 16-bit Origin ASNs has perhaps been unavoidable, it is also somewhat saddening. The number of 32-bit-only Origin ASNs increased across all categories—IPv4＋IPv6, IPv4 only, and IPv6—though the increases were smaller than two years ago. Notably, IPv4 only still only saw a three-digit rise that was smaller than the decline in 16-bit AS numbers, resulting in a decline in total IPv4 only Origin ASNs for a second year running. Another trend, ongoing since 2021, is that the increase in IPv4＋IPv6 32-bit only ASes has been exceeding the increase in IPv4 only ASes, and as such, we will be watching next year's results with interest.

Table 3: Total Number of Unique IPv4 Addresses in Full IPv4 Routes and Total Number of Unique IPv6 /64 Blocks in Full IPv6 Routes

Date	No. of IPv4 addresses	No. of IPv6 /64 blocks
Sep. 2015	2,791,345,920	31,850,122,325
Sep. 2016	2,824,538,880	26,432,856,889
Sep. 2017	2,852,547,328	64,637,990,711
Sep. 2018	2,855,087,616	258,467,083,995
Sep. 2019	2,834,175,488	343,997,218,383
Sep. 2020	2,850,284,544	439,850,692,844
Sep. 2021	3,036,707,072	461,117,856,035
Sep. 2022	3,068,374,784	532,578,391,219
Sep. 2023	3,055,604,992	700,359,397,494
Sep. 2024	3,033,333,504	896,502,953,452

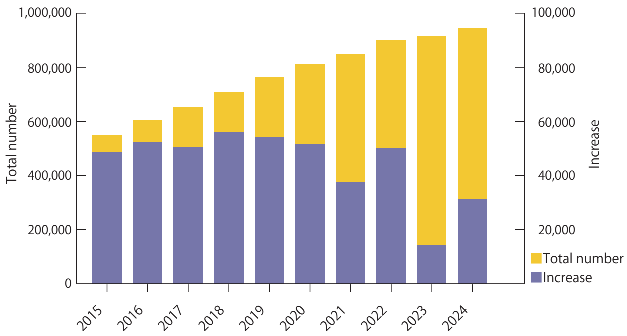

Figure 1: Total Number of Full IPv4 Routes and Annual Increases

Table 4: IPv4/IPv6 Full-Route Origin AS Numbers

ASN	16-bit(1～64495)					32-bit only(131072～4199999999)				
Advertised route	IPv4+IPv6	IPv4 only	IPv6 only	total	(IPv6 -enabled)	IPv4+IPv6	IPv4 only	IPv6 only	total	(IPv6 -enabled)
Sep. 2015	8228	34544	137	42909	(19.5%)	1424	6801	78	8303	(18.1%)
Sep. 2016	9116	33555	158	42829	(21.7%)	2406	9391	146	11943	(21.4%)
Sep. 2017	9603	32731	181	42515	(23.0%)	3214	12379	207	15800	(21.7%)
Sep. 2018	10199	31960	176	42335	(24.5%)	4379	14874	308	19561	(24.0%)
Sep. 2019	10642	31164	206	42012	(25.8%)	5790	17409	432	23631	(26.3%)
Sep. 2020	11107	30374	229	41710	(27.2%)	7653	19668	574	27895	(29.5%)
Sep. 2021	11465	29219	302	40986	(28.7%)	9514	21108	5242	35864	(41.1%)
Sep. 2022	11613	28398	369	40380	(29.7%)	10816	22211	5764	38791	(42.7%)
Sep. 2023	11770	27617	460	39847	(30.7%)	12640	22128	2067	36835	(39.9%)
Sep. 2024	12068	26720	476	39264	(31.9%)	13905	22737	2386	39028	(41.7%)

Topic 2

DNS Query Analysis

IIJ provides a full-service resolver to enable DNS name resolution for its users. Here, we discuss the state of name resolution, and analyze and reflect upon data from servers provided mainly for consumer services, based on a day's worth of full-service resolver observational data obtained on October 9, 2024.

The full-service resolver provides a name resolution function that replies to DNS queries from user devices. Specifically, to resolve a name, it starts by looking at the IP address of an authoritative server for the root zone (the highest level zone), which it queries, and then goes through other authoritative servers to find the records it needs. If the full-service resolver repeatedly queries other servers like this, it can result in increased load and delays, so the information obtained is cached, and when the same query is received again, the response is sent from the cache. Recently, DNS-related functions are implemented on devices that lie on route paths, such as consumer-level routers and firewalls, and these devices are sometimes also involved in relaying DNS queries and applying control policies. Some applications, such as Web browsers, also have their own implementations of name resolver functionality and in some cases resolve names based on a policy that differs from the OS settings.

ISPs notify users of the IP address of full-service resolvers via various protocols, including PPP, DHCP, RA, and PCO, depending on the connection type, and they enable automatic configuration of which nameserver to use for name resolution on user devices. ISPs can notify users of multiple full-service resolvers, and users can specify which nameserver to use by altering settings in their OS, browser, or elsewhere. When more than one nameserver is configured on a device, which one ends up being used depends on the device's implementation or the application, so any given full-service resolver is not aware of how many queries a user is sending in total. When running full-service resolvers, therefore, this means that you need to keep track of query trends and always try to keep some processing power in reserve because changes in behavior or status on the user end can conceivably result in a sudden increase in queries to a given resolver.

Observational data on the full-service resolver provided by IIJ show fluctuations in user query volume throughout the day, with volume hitting a daily trough of about 0.15 queries/sec per source IP address at around 4:25 a.m., and a peak of about 0.32 queries/sec per source IP address at around 9:50 p.m. Overall volume was down 0.02pt vs. the previous year. The breakdown shows that IPv4 accounted for around 59% of queries and IPv6 for around 41%, with IPv6's share having risen by around 1pt from the previous year.

Recent years have seen a tendency for queries to rise briefly at certain round-number times, such as on the hour marks in the morning. As in the past, the number of queries increased around 6 a.m. and 7 a.m. We observed increases in query volume at 14 and 9 seconds before the 6 a.m. and 7 a.m. marks, as was observed in the previous year. This is a pattern we have seen in recent years, with query volume rising sharply at the hour mark and then tapering off gradually, but with the sudden spikes that occur ahead of the hour mark, query volume quickly returns to roughly where it had been. Hence, because a large number of devices are sending queries in almost perfect sync, we surmise that lightweight, quickly completed tasks of some sort are being executed. This year, however, we noticed that these increases at round-number times were smaller than in the past, and we also observed a tendency for query volume over the period from 8 a.m. to 10 p.m. to actually fall at the top of each hour and then gradually increase. We suspect this reflects some implementation changes on client devices that use name resolution.

Turning to protocols, UDP accounted for almost all (98.438%) of the queries. That said, TCP queries have been rising over the last few years, from 0.189% of total in 2021 to 0.812% in 2022, 1.419% in 2023, and 1.561% in 2024. Possibly the main driver of this is an increase in queries using DNS over TLS (DoT). DoT basically uses TCP port 853 to send queries, so an increase in the use of DoT results in an increase in TCP queries.

Looking at the query record types, A records that query the IPv4 address corresponding to the host name, AAAA records that query IPv6 addresses, and HTTPS records used to resolve Web services account for 98% of the total. The trends in A and AAAA queries differ by IP protocol, with more AAAA record queries being seen for IPv6-based queries. Of IPv4- based queries, around 62% are A record queries and 17% AAAA record queries (Figure 2). With IPv6-based queries, meanwhile, A record queries account for around 40% and AAAA record queries around 35% of the total (Figure 3). Compared with the previous year, we observe a 5-percentage-point increase in A record queries for IPv4 and a 2-percentage-point increase for IPv6. Meanwhile, HTTPS record queries, which we started to see in 2020, declined for the first time. They accounted for around 17% of IPv4 and 24% of IPv6 queries, decreases of 3 percentage points for IPv4 and 2 percentage points for IPv6 from the previous year. This may be due to some sort of changes in client implementations. SVCB records, which we started to see in 2022, accounted for 0.30% of IPv4 and 0.57% of IPv6 queries, and while these queries are still only a small fraction of the total, they are progressing steadily. This may be attributable to the use of implementations of Discovery of Designated Resolvers (DDR), a newly proposed protocol designed to allow clients to detect encryption-capable full-service resolvers.

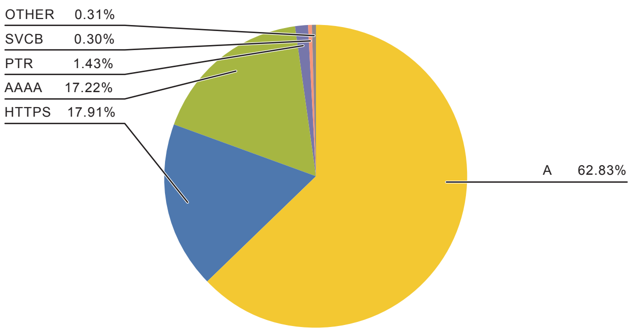

Figure 2: IPv4-based Queries from Clients

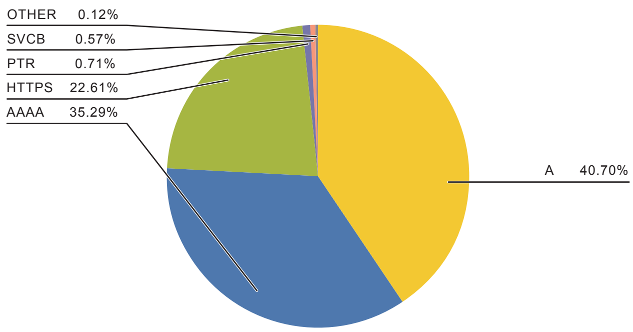

Figure 3: IPv6-based Queries from Clients

Topic 3

IPv6 & Mobile

In this section, we again report on the volume of IPv6 traffic on the IIJ backbone, source ASNs, and the main protocols used. We also go over the state of IPv6 connections on mobile services by device OS.

■ Traffic

Figure 4 shows traffic measured using IIJ backbone routers at core POPs (points of presence—3 in Tokyo, 2 in Osaka, 2 in Nagoya). The data cover the eight months from February 1 to September 30, 2024. Both IPv6 and IPv4 Internet traffic volumes saw a moderate downtrend during that period. Both IPv6 and IPv4 traffic were up when viewed alongside figures for the same day of the previous year (lighter lines on the graph), with IPv6 traffic rising 14.309% and IPv4 traffic rising 14.505% year over year, so almost the same rate of increase for each.

Figure 5 graphs traffic indexed to 100 as of February 1, 2024. As noted, traffic volumes declined mildly from the start of the year, with roughly similar moves for both IPv6 and IPv4.

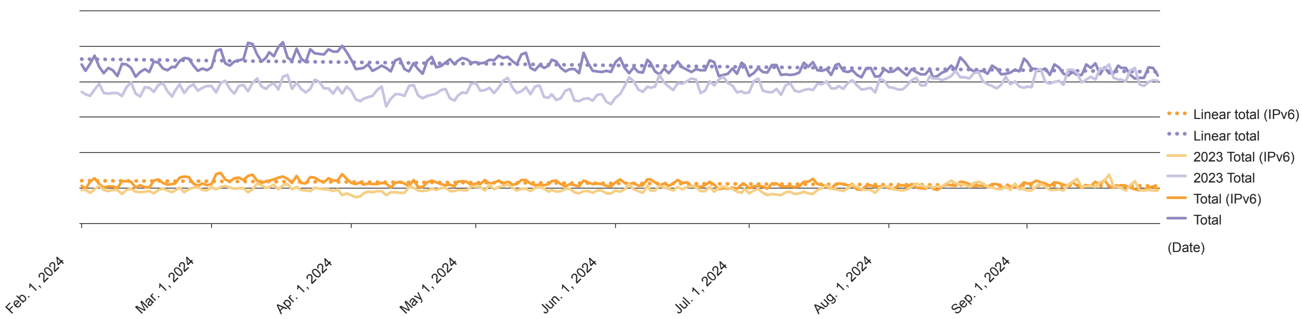

Figure 4: Traffic Measured on Backbone Routes at IIJ's Core POPs

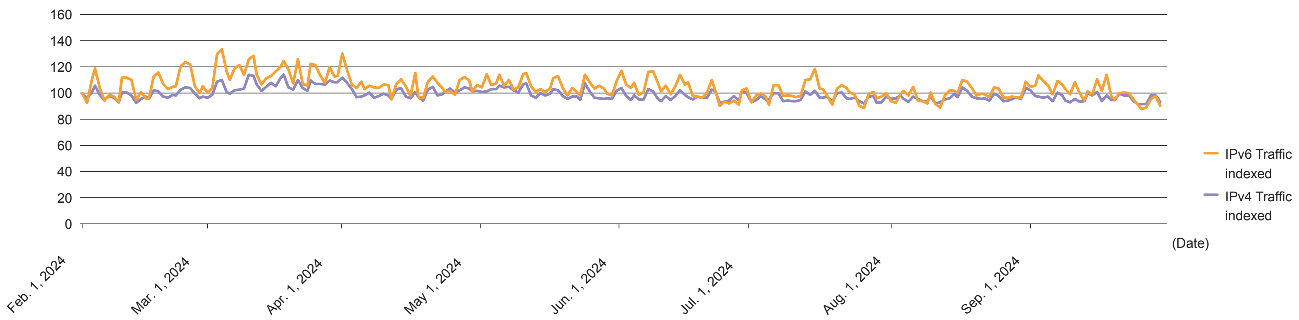

Figure 5: Traffic Indexed to 100 as of February 1

Next, Figure 6 shows IPv6 as a proportion of total traffic. This moves between a minimum of 18.6% and a maximum of 22.5%, averaging 20.16% during the observation period. No major trends are discernible, and the figures are largely in line with those from a year earlier; IPv6 traffic growth perhaps experienced a minor lull.

Table 5 tracks the IPv6 ratio since 2017. It has grown steadily with the exclusion of the Covid period, with this year's ratio being on par with last year's.

■ Traffic Source Organization (BGP AS)

Next, Figures 7 and 8 show the top IPv6 and IPv4 traffic source organizations (BGP Source AS Number) for February 1 – September 30, 2024.

For IPv6, traffic within IIJ accounts for over 60% of the total. Looking at non-IIJ ASes, Company A, a major US search provider, moved up from No. 2 last year to take the top spot with 6% of traffic. At No. 2 with 5% was Company B, a major Japanese content provider, which was in No. 1 last year. While No. 1 and 2 switched places this year, the traffic volumes are neck and neck, so they may well trade places again in the future. At No. 3 is Company C, a US cloud operator, up from No. 8 last year. Having been at No. 16 two years ago, it has increased its IPv6 traffic substantially over the past few years. But at around 2% of total traffic, it is a fair way behind the top two companies, and with volumes for No. 4 and below also being very close, we can expect to see the rankings reshuffled on a yearly basis going forward as well.

Table 5: IPv6 as a Proportion of Total Traffic (Since 2017)

	2017 IIR Vol. 37	2018 IIR Vol.41	2019 IIR Vol.45	2020 IIR Vol.49	2021 IIR Vol.53	2022 IIR Vol.57	2023 IIR Vol.61	2024 IIR Vol.65
IPv6 ratio	4%	6%	10%	10%	11.2%	15.1%	20.1%	20.16%

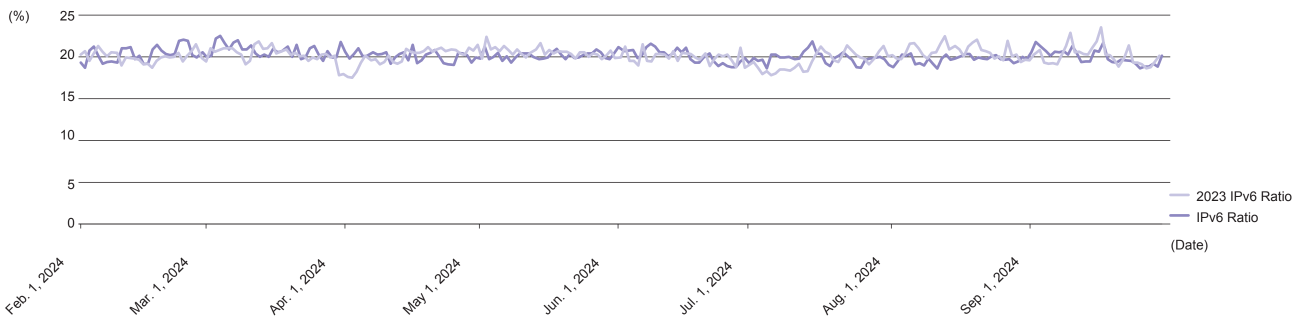

Figure 6: IPv6 as a Proportion of Total Traffic

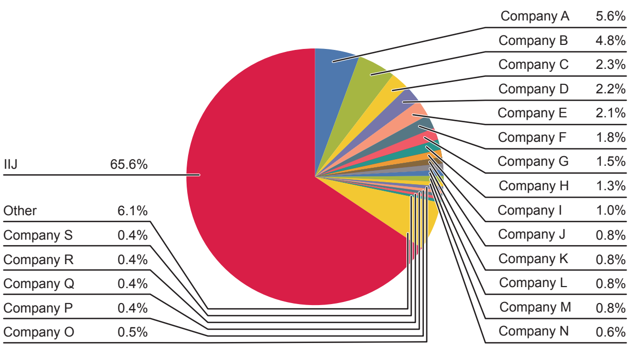

Figure 7: IPv6 Traffic by Source Organization (BGP AS Number)

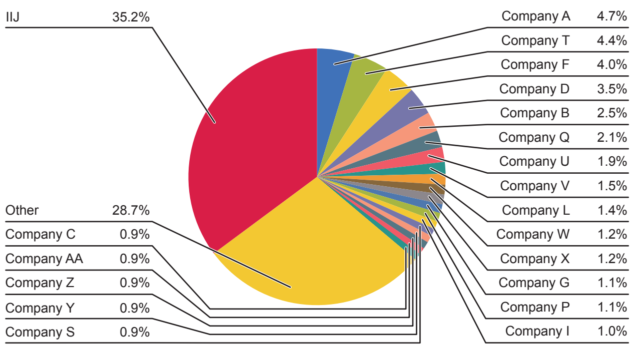

Figure 8: IPv4 Traffic by Source Organization (BGP AS Number)

Protocols Used

Figure 9 plots IPv6 traffic according to protocol number (Next Header) and source port number, and Figure 10 plots IPv4 traffic according to protocol number and source port number (for the week of Monday, September 30 – Sunday, October 6, 2024).

In the IPv6 space, similar to last year, the top four protocols—HTTPS, QUIC, NAT Traversal, and ESP in that order—accounted for 91% of usage. HTTPS accounted for 74% and QUIC 9%, so HTTP-related protocols accounted for over 80%, while VPN-related protocols NAT-T and ESP accounted for 8.4%.

While traffic patterns remain largely similar to last year, there appears to be an overall increase in daytime traffic. IPv6 traffic, in particular, increased during the day on Saturdays and Sundays, suggesting that individual user traffic may constitute a significant portion of the total. In the IPv4 space, UDP443 has now surpassed TCP80, indicating a decline in unencrypted HTTP traffic. Yet while IPv6 HTTP traffic is almost imperceptible on the graph, IPv4 HTTP traffic has only just now been overtaken by QUIC, suggesting that many old servers are still in operation.

IPv6 on Mobile Devices

Following on from our installment of this report in IIR Vol. 57 two years ago (https://www.iij.ad.jp/en/dev/iir/057.html) and in last year's IIR Vol. 61 (https://www.iij.ad.jp/en/dev/iir/061.html), we again look at IPv6-enabled rates on personal IIJmio Mobile Service connections. We also look at differences by device OS and at whether there are differences depending on device manufacturer.

The IPv6-enabled rate for devices connected to the IIJmio Mobile Service was 60.6%. This represents an annual increase of around 2 percentage points, from 58.73% last year and 56.3% the year before that. By device OS, 87.010% of Apple iOS devices had IPv6 enabled, while the figure was 30.235% for Android devices. The Android

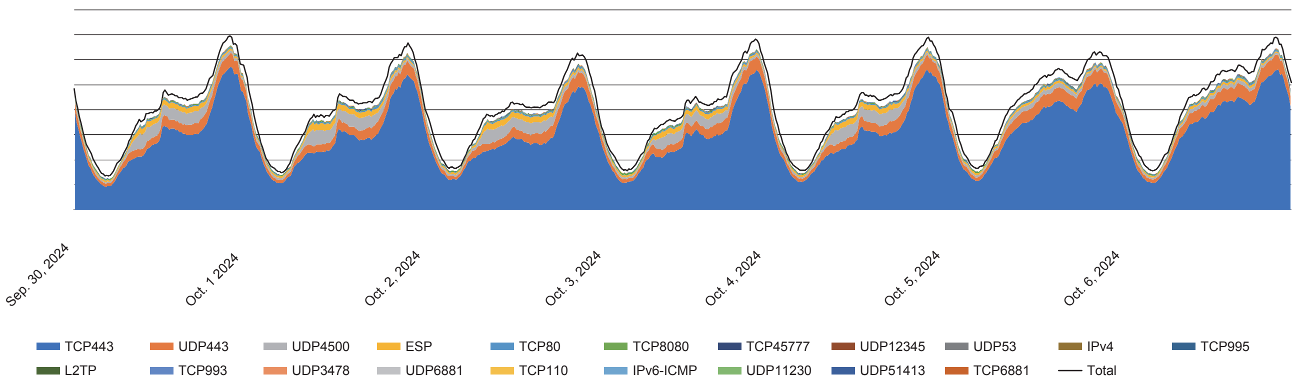

Figure 9: Breakdown of IPv6 Traffic by Source Port Number

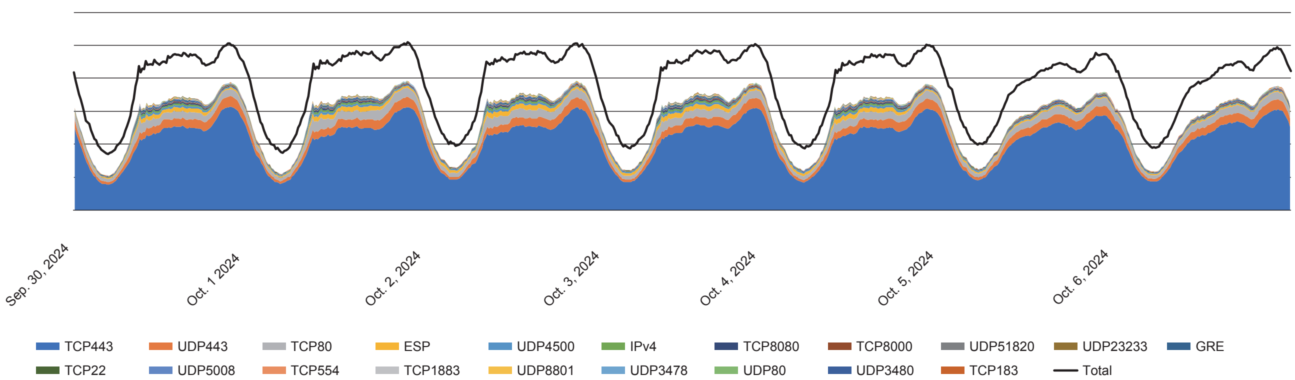

Figure 10: Breakdown of IPv4 Traffic by Source Port Number

IPv6-enabled rate was up a substantial 5 percentage points from last year, which contributed to the overall rise in the IPv6-enabled rate.

Next, we look at IPv6-enabled rates by manufacturer for the top 20 devices connected to the IIJmio Mobile Service. Figure 11 graphs the top 20 spots. Apple products are widely used in Japan and thus account for over 53%. Apple's IPv6-enabled rate was high at around 87%, a slight increase from last year. The manufacturer with the highest IPv6-enabled rate was Motorola Mobility at 91.7%, followed by Apple at 87% and Google in third place at 86%.

Here, we highlight FCNT, which ranks 14th by number of IIJmio-connected devices. FCNT released new products such as the arrows We2 this year, and arrows We2 alone had an IPv6-enabled rate of 97.6%. This is likely because its default APN profile is set to PDP-Type IPv4v6, enabling IPv6 connectivity. Yet the F-51B, also from FCNT and designed for MNOs, had an IPv6-enabled rate of just 7.4%, and so the overall IPv6-enabled rate for FCNT devices was only 31.2%.

■ Summary

We have examined traffic on the IIJ Internet backbone core, source ASNs, and main protocols used. Traffic volume declined slightly during the observation period but was up by over 14% vs. last year. IPv6 usage remained largely flat at 20.16%, essentially treading water this year. Looking at IPv6 traffic by source AS, No. 1 and 2 traded places but were still neck and neck, while No. 3 and below also remained closely matched at around 1–2% of total traffic each.

As has been the case so far, IPv6 is more widely used on relatively new servers, with encrypted HTTP protocols accounting for over 80% of traffic and unencrypted HTTP only a sliver. As for IPv4 traffic, however, unencrypted HTTP is still fairly prevalent.

In the mobile space, Android OS devices saw a 5-percentage-point increase in the IPv6-enabled rate, while the overall increase was just under 2 percentage points. We are also seeing new devices released with IPv6 enabled via the PDP-Type setting in their APN configuration. It would be good to see the number of devices with IPv6 enabled by default continue to rise ahead.

We will continue to watch the IPv6 situation from a range of angles and provide updates as new developments come to light.

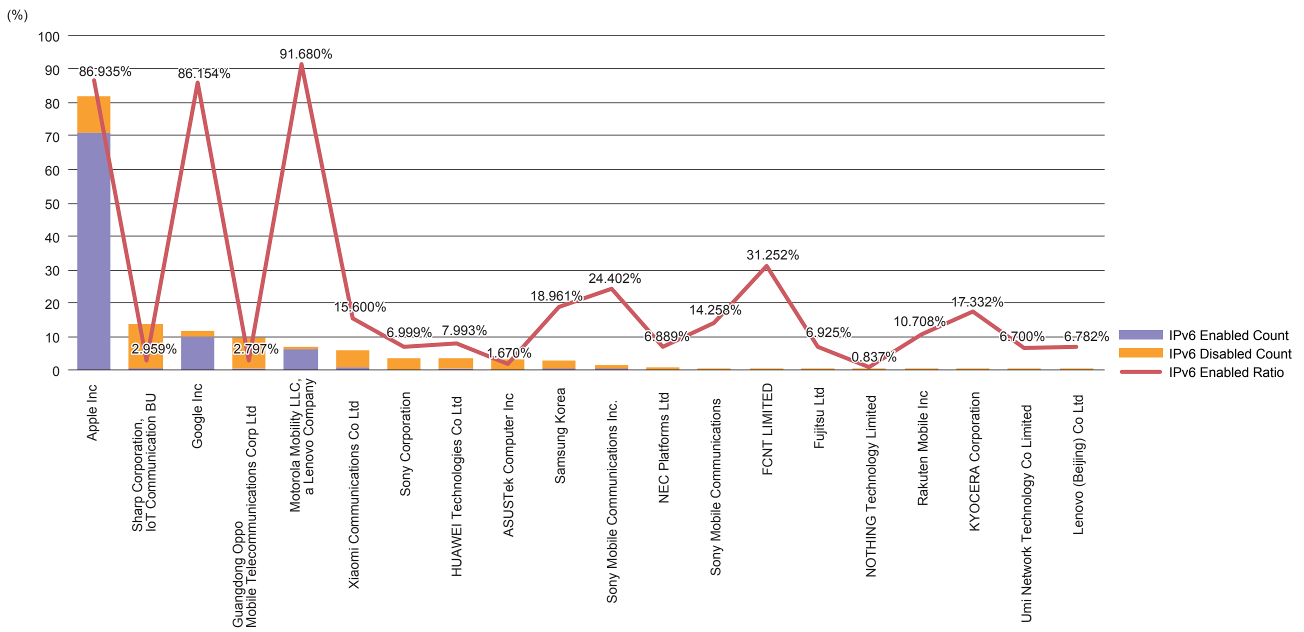

Figure 11: IPv6 Support by UE Manufacturer (Top 20)

Topic 4

Internet Backbone Trends

In this section, we cover recent Internet backbone trends from the perspective of IIJ's internet backbone infrastructure interconnectivity and routing.

Interconnection Interface Trends and Requirements

To facilitate interconnectivity on the Internet, service operators need to adopt standardized interfaces. Currently, IIJ primarily uses 400G-FR4, 100G-LR4, and 10G-LR for interconnection interfaces. There has been a trend in recent years of reviewing interconnectivity between operators that is based on 10G interfaces. The main objectives for establishing interconnectivity are to reduce traffic exchange costs and to improve communication quality through increased efficiency. From a cost reduction perspective, using 10G interfaces for interconnectivity is perhaps becoming less cost-effective. The main reasons for this are improving cost per 100G port on interconnection routers used at the AS boundaries between operators and the relative difficulty of managing 10G interfaces. At IIJ, we also request/select interconnection routers offering both high bandwidth and high port density. Using 10G interfaces on such routers results in lower port utilization efficiency. This is because, in terms of port configuration, the use of 100G and 10G ports is often mutually exclusive. Physical interface ports are a scarce resource, and when they are used for 10G instead of 100G connections, bandwidth is lower and thus less traffic can be carried, which results in reduced port utilization efficiency. And trying to select routers with interface support covering 10G all the way through to 400G for interconnection applications limits the range of options that are satisfactory in terms of functionality and price compared with routers that support only 100G/400G.

At IIJ, therefore, we ask interconnection partners to upgrade from 10G to 100G interfaces when 10G interfaces are not absolutely necessary, when multiple 10G lines are already bundled through link aggregation, or when bandwidth might be increased beyond 10G in the future. Similarly, IIJ also receives requests from its interconnection partners to either switch from individual 10G connections to IXP-based connections or to upgrade ports to 100G interfaces.

By comparing Figures 12 and 13, we can see how the percentage breakdown of interconnection interfaces on IIJ's internet backbone has changed over the past year. Over this period, the use of 10G interfaces declined as a percentage total while the use of 100G interfaces increased.

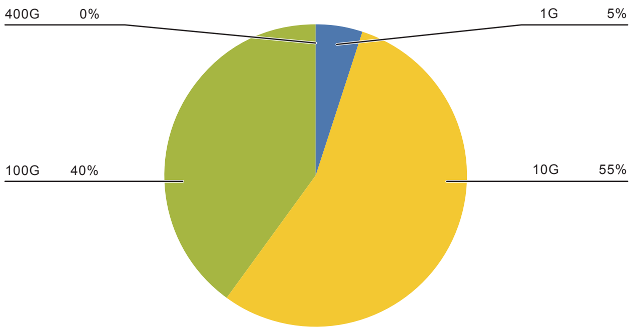

Figure 12:Breakdown of Interconnection Interfaces on IIJ's Internet Backbone (October 2023)

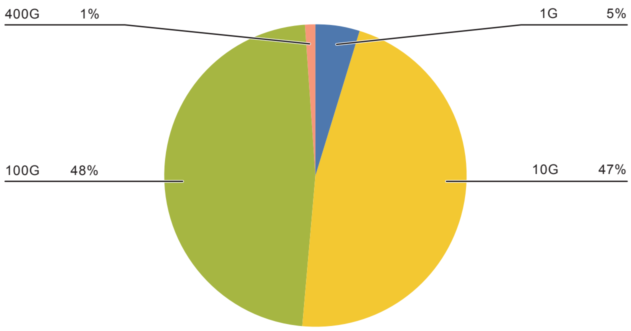

Figure 13:Breakdown of Interconnection Interfaces on IIJ's Internet Backbone (October 2024)

While 400G interfaces are being considered as an interconnection interface requirement for connections of 100G and above, adoption still appears limited. While IXP operators seem to be making relatively decent progress in supporting 400G, it likely remains more of a future consideration for interconnections between service providers, except those with large traffic volumes.

Another new trend we are seeing is efforts to promote the use of 100G-LR interfaces within the 100G category. Compared with the currently widespread 100G-LR4, the 100G-LR specification increases the transmission capacity per wavelength from 25G to 100G. Advantages of this include lower cost per unit due to the number of lasers per unit being reduced from four to one and reduced failure rates due to there being fewer components. The use of 100G-LR on interconnection interfaces has begun primarily at IXPs in Europe and North America, and it is expected to be taken up on interconnections between service providers ahead.

Current State of RPKI

Here, we provide an update on the state of RPKI. We look at current data on ROAs, signed objects that verify which IP addresses are owned by an organization. ROA registration data from the NIST RPKI Monitor (Figures 14 and 15), available online*1, show that among all IPv4 routes on the Internet, 53.38% are Valid (ROAs registered, routing verified), 46.15% are Not-Found (ROA not yet registered), and 0.47% are Invalid (discrepancies in ROA registration, treated as unauthorized routes). In the IPv6 space, 55.30% of routes are Valid, 40.24% are

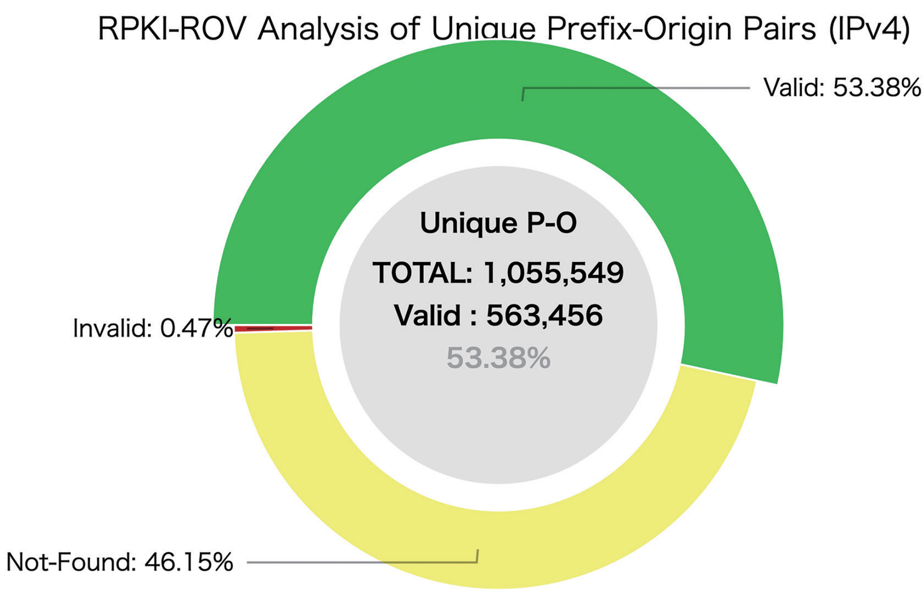

Figure 14:ROA Registration Data from the RPKI Monitor (IPv4) as of 18:00, October 15, 2024

*1 See NIST RPKI Monitor (https://rpki-monitor.antd.nist.gov/).

Not-Found, and 4.46% are Invalid. With the ROA registration rate for IPv4 having surpassed 50%, legitimacy has been established for more than half of all routes, which suggests substantial progress has been made. Meanwhile, over 4% of IPv6 routes are Invalid. Registering an ROA is simple: you register the Origin AS that will generate and advertise the route, prefix and subnet information, and the maximum route length (subnet size). Any of the parameters being inconsistent with the ROA will result in the route being marked as Invalid, so such discrepancies need to be corrected unless they are for testing purposes.

Currently, we have enabled RPKI-ROV for AS boundary BGP peers and perform ROA-based origin validation on routes that peers advertise to IIJ. Since Invalid routes cannot be distinguished from route hijacking attempts, our policy is basically to reject them.

Now let's look at the ROA registration status of Internet routes that IIJ generates and advertises (Table 6). IIJ participates in the Internet using Global AS number AS2497, and thus the Origin AS for routes IIJ advertises is AS2497. As of October 16, 2024, 44.0% of routes originating from AS2497 are Valid.

	No. of routes originated by AS2497	No. of routes IIJ transits and advertises to the Internet
Valid	80	4200
Unknown	102	3757
Invalid	0	3
Valid rate	44.0%	52.8%

Table 6: ROA Registration Data from the RPKI Monitor (IPv6) as of 18:00, October 15, 2024

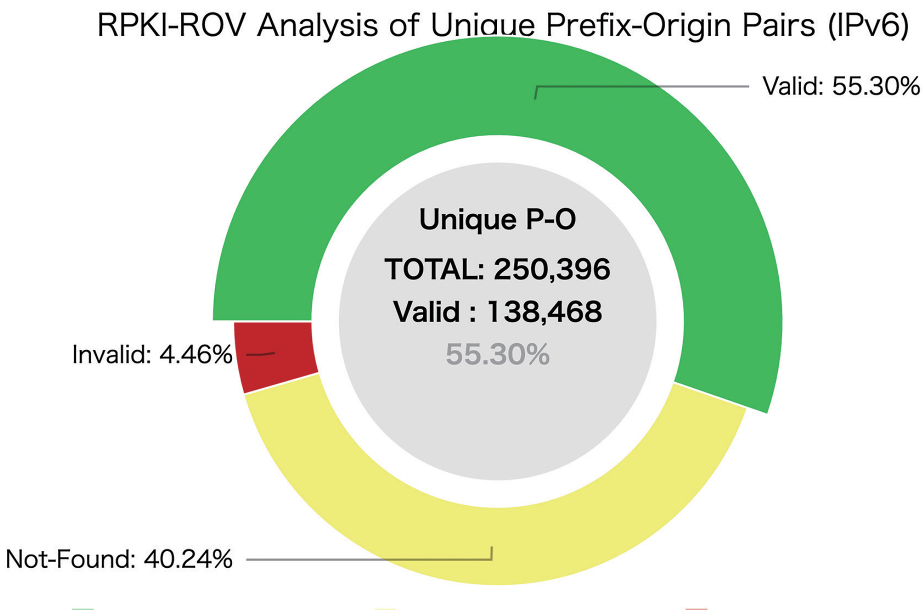

Figure 15: ROA Registration Data from the RPKI Monitor (IPv6) as of 18:00, October 15, 2024

This means ROAs are registered, confirming legitimacy, for just under half. Yet, ROV returns Not-Found for over 50% of the routes, indicating no ROA has been registered. One significant obstacle to ROA registration is the need for organizations that own IP addresses to register and issue their own ROAs. While IIJ has registered ROAs for nearly all IP addresses it owns (excluding special cases), there is evidently a lack of progress among users who bring their own IP addresses to IIJ. Since IIJ cannot currently register ROAs on behalf of users, they must complete the registration themselves. IIJ offers support in this regard, so we encourage you to take advantage of this to help us increase the ROA registration rate for AS2497.

When including routes for which IIJ provides transit service to customers, 52.8% of all routes are Valid according to ROV. This indicates decent progress in terms of ROA issuance within Japan as well.

1.BGP and Routes
Tomohiko Kurahashi
Infrastructure Engineering Division Operation Engineering Department Operation Research & Development Section, IIJ

2.DNS Query Analysis
Yoshinobu Matsuzaki
Infrastructure Engineering Division Operation Engineering Department Operation Research & Development Section, IIJ

3.IPv6 & Mobile
Taisuke Sasaki
MVNO Division Infrastructure Development Department, IIJ

4.Internet Backbone Trends
Yuichi Yomogita
Infrastructure Engineering Division Network Engineering Department Network Engineering Section 1, IIJ