Internet Trends as Seen from IIJ Infrastructure —2020

Internet services provider IIJ operates some of the largest network and server infrastructure in Japan. Here, we report on Internet trends over the past year based on information obtained through the operation of this infrastructure. We analyze changes in trends from the perspective of BGP routes, DNS query analysis, IPv6, and mobile. We also discuss conditions observed before the deployment of BGP ROV on the IIJ backbone.

Topic 1

BGP 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 2). During the past year, RIPE NCC and LACNIC completely exhausted their IPv4 address pools. Only APNIC and AfriNIC remain, but APNIC has already projected that it will run out at the beginning of 2021, and in January 2020 AfriNIC placed a size limit (/22) on allocations/assignments.

The increase in the number of routes per year has been trending downward since peaking in 2018, but the total number now exceeds 800,000. The /22 prefix now accounts for over 100,000 routes, but the proportion of routes accounted for by the /22, /23, and /24 prefixes rose only slightly to 80.9% of the total.Meanwhile, the number of unique IPv4 addresses, although increasing vs. 2019, when it fell for the first time since 2011, remains below its 2018 and 2017 levels. Whether 2018 will turn out to have been the peak here as well will bear watching ahead.

Next we take a look at IPv6 full-route data (Table 3). In November 2019, ARIN received its first additional /12 block allocation (the second such allocation since RIPE NCC received one in June 2019).

The total number of routes increased by about the same amount as in the previous year and now exceeds 90,000. We expect it to surpass 100,000 in 2021. Also, 50.0% of all routes and 58.1% of those constituting the increase are

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. 2011	19	12	27	81	233	457	794	1407	11909	5907	9885	19515	26476	26588	35515	34061	190276	363162
Sep. 2012	19	14	29	84	236	471	838	1526	12334	6349	10710	20927	30049	31793	42007	39517	219343	416246
Sep. 2013	16	11	30	93	250	480	903	1613	12748	6652	10971	22588	32202	34900	48915	42440	244822	459634
Sep. 2014	16	12	30	90	261	500	983	1702	13009	7013	11659	24527	35175	37560	54065	47372	268660	502634
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

Table 2: Total Number of Unique IPv4 Addresses in Full IPv4 Routes

Date	No. of IPv4 addresses
Sep. 2011	2,470,856,448
Sep. 2012	2,588,775,936
Sep. 2013	2,638,256,384
Sep. 2014	2,705,751,040
Sep. 2015	2,791,345,920
Sep. 2016	2,824,538,880
Sep. 2017	2,852,547,328
Sep. 2018	2,855,087,616
Sep. 2019	2,834,175,488
Sep. 2020	2,850,284,544

/48 routes, from which we infer that IPv6 is also rolling out smoothly on end sites.

Lastly, let's also take a look at IPv4/IPv6 full-route Origin AS figures (Table 4). In the past year, an additional 3072 32-bit-only AS numbers were allocated to RIPE NCC and 2048 to LACNIC.

Both the decrease in 16-bit Origin Autonomous System Numbers (ASNs) and the increase in 32-bit-only Origin ASNs were around the same as those in the previous year, and 32-bit-only ASNs now account for 40% of all Origin ASNs. IPv6-enabled ASNs, which advertise IPv6 routes, continue to rise and now account for 28.1% of the total. Whether this will far exceed 30% or not in 2021 will be a point to watch.

Topic 2

DNS Query Analysis

IIJ provides a full resolver to enable DNS name resolution for its users. In this section, 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 resolver observational data obtained on September 30, 2020.

The full resolver starts by looking at the IP address of an authoritative name server for the root zone (the highest level zone), which it queries, and then goes through other authoritative nameservers to find the records it needs. Queries repeatedly sent to the full resolver 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 broadband routers and firewalls, and these devices are sometimes 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 without relying on OS settings.

Table 3: 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. 2011	68	13	22	3530	406	248	45	87	95	2356	6870
Sep. 2012	102	45	34	4448	757	445	103	246	168	3706	10054
Sep. 2013	117	256	92	5249	1067	660	119	474	266	5442	13742
Sep. 2014	134	481	133	6025	1447	825	248	709	592	7949	18543
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

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. 2011	4258	32756	115	37129	(11.8%)	90	1278	13	1381	(7.5%)
Sep. 2012	5467	33434	125	39026	(14.3%)	264	2565	17	2846	(9.9%)
Sep. 2013	6579	34108	131	40818	(16.4%)	496	3390	28	3914	(13.4%)
Sep. 2014	7405	34555	128	42088	(17.9%)	868	4749	55	5672	(16.3%)
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%)

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

Observational data on the full resolver provided by IIJ show fluctuations in user query volume throughout the day, with volume hitting a daily trough of about 0.06 queries/sec per source IP address at around 4:30 a.m., and a peak of about 0.24 queries/sec per source IP address at around 12:30 p.m. These values are almost the same as last year, with a slight increase of 0.01 points from early morning into the daytime.

Broken down by protocol (IPv4 and IPv6), IPv4 queries per IP address fell vs. the previous year. While there is no major change in the middle of the night, we observe up to a 0.03-point drop in the daytime and from evening through nighttime. Meanwhile, IPv6 queries per IP address are up by around 0.03 points across all times of day, including late night. This suggests that IPv6-capable devices are gradually making their way into the home and that existing devices are being replaced. And looking at total query count, both the number of source IPs and the number of actual queries are higher for IPv6 than for IPv4. The number of IPv6-based queries is on the rise, accounting for around 63% of the total, up by 3 points from 60% in 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. The number of query sources also increases, with a particularly noticeable pattern around 7 a.m., which is possibly due to tasks scheduled on user devices and increases in automated network access that occur when devices are activated by, for example, an alarm clock function. In the previous year, we noted increases in queries 14 seconds and 10 seconds before every hour mark, and the 2020 results also show another increase 20 seconds before every hour. The increase in queries that occurs on the hour tapers off gradually, but with the spikes that occur before 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.

For example, there are mechanisms for completing basic tasks, such as connectivity tests or time synchronization, before bringing a device fully out of sleep mode, and we posit that the queries used for these tasks are behind the spikes.

Looking at the query record types, most are A records that query the IPv4 address corresponding to the host name and AAAA records that query IPv6 addresses. 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 79% are A record queries and 15% AAAA record queries (Figure 1). With IPv6-based queries,

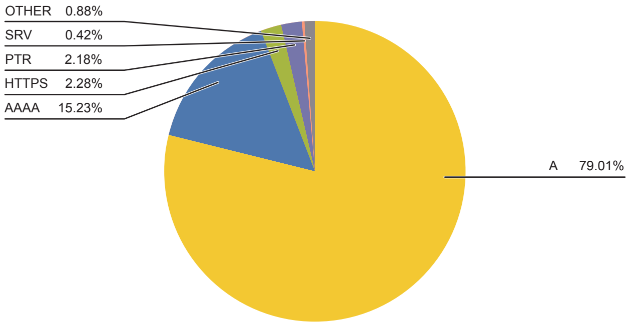

Figure 1: IPv4-based Queries from Clients

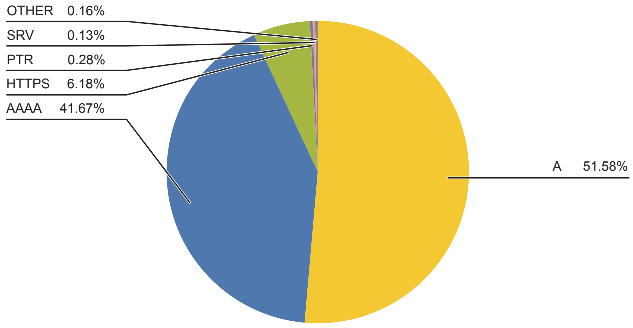

Figure 2: IPv6-based Queries from Clients

meanwhile, AAAA record queries account for a higher share of the total, with around 51% being A record and 41% being AAAA record queries (Figure 2). Compared with the previous year, we observe drops in A record queries of 5 percentage points for IPv4 and 3 percentage points for IPv6. HTTPS-type records, newly implemented in 2020, accounted for some 2% of IPv4 and 6% of IPv6 queries, coming in behind the A and AAAA query volumes. Systems currently supporting DNS over HTTPS include Apple's iOS 14, and we expect these queries to rise gradually as implementations spread.

Topic 3

IPv6

In this section, we report on the volume of IPv6 traffic on the IIJ backbone, source ASNs, and the main protocols used.

Traffic

As before, we again present IPv4 and IPv6 traffic measured using IIJ backbone routers at core POPs (points of presence—Tokyo, Osaka, Nagoya), shown in Figure 3. The data span the year from October 1, 2019 to September 30, 2020.

Traffic trends in 2020 differed from what we observed up till the previous year, with COVID-19 being a factor from the year's outset. Although no major changes were apparent until about February, IPv4 traffic increased substantially from March as COVID-19 prompted school closures and Japan's state of emergency declaration, resulting in people staying home. As discussed in Vol. 48 (https://www.iij.ad.jp/en/dev/iir/048.html), mobile traffic fell during the stay-at-home period, and fixed broadband services and corporate VPN services saw increases, which is likely behind the increase in IPv4 traffic.

To see the relative increases and decreases during the observation period, we graphed normalized IPv4 and IPv6 traffic with the values for the first day (October 1, 2019) indexed to 1 (Figure 4). The middle of the graph around April and May corresponds to Japan's state of emergency, when people were staying home. IPv4 traffic was up about 8% during the period, while IPv6 traffic looks to have fallen. Once the state of emergency ended, the increase in IPv4 traffic settled down to a level representing a slight increase vs. the start of the observation period. IPv6 traffic also ultimately saw a mild increase.

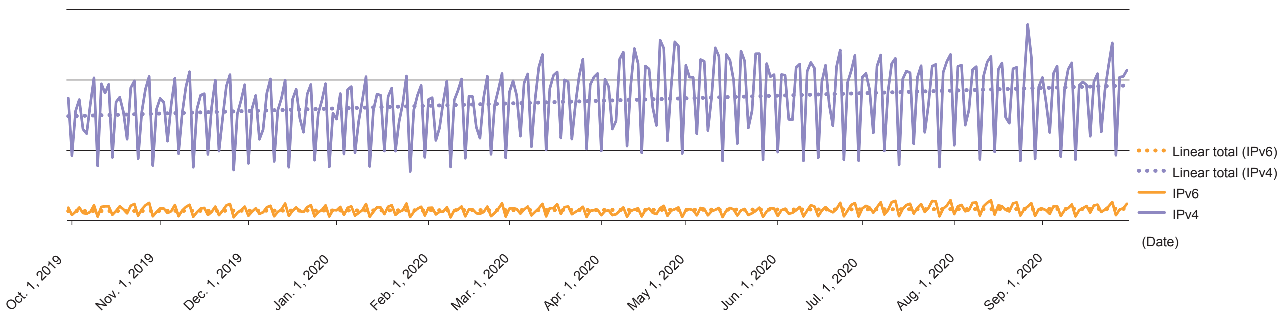

Figure 3: IPv4 Traffic and IPv6 Traffic (Oct. 2019 – Sep. 2020)

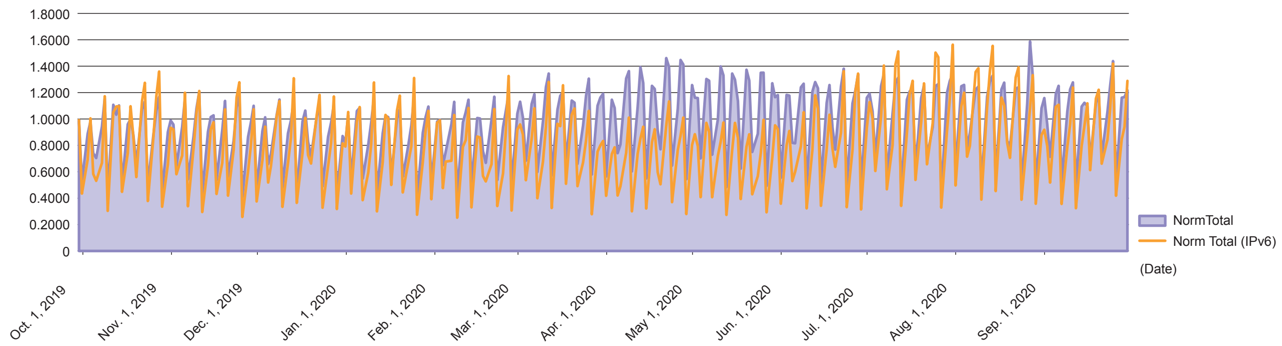

Figure 4: IPv4 Traffic and IPv6 Traffic, Indexed to Initial Value of 1

As Figure 5 shows, IPv6 declined as a percentage of the total during the state of emergency, but it eventually returned to about its initial level.

■ Traffic Source Organization (BGP AS)

Next, Figures 6 and 7 show the top annual average IPv6 and IPv4 traffic source organizations (BGP AS Number) for the year from October 2019 through September 2020.

Company A retains the top spot, but its share of traffic is down 5 percentage points since last time we reported. Traffic with IIJ's ASN as the source has grown substantially, and while this could be due to peculiarities of the observation point, the main factor is probably the growth in IPv6 video streaming traffic on JOCDN's platform, like last year. We observed no other major shifts or noticeable trends.

■ Protocols Used

Figure 8 plots IPv6 traffic according to protocol number (Next Header) and source port number, and Figure 9 plots IPv4 traffic according to protocol number and source port number (for the week starting Monday, October 5, 2020).

In the IPv6 space, TCP 443 (HTTPS) came in at No. 1 and UDP 43 (QUIC) at No. 2, so over 80% was attributable to HTTP encryption protocols. Of particular note this time around, TCP 80 (unencrypted HTTP) fell to No. 4 and ESP (IPSec encryption) came in at No. 3. ESP is observed more during the daytime on weekdays and is scarce on weekends, which probably indicates an increase in IPv6 VPN usage on corporate networks. It accounts for some 19% of traffic during its most prominent period around noon on weekdays, and even eclipses QUIC during the daytime.

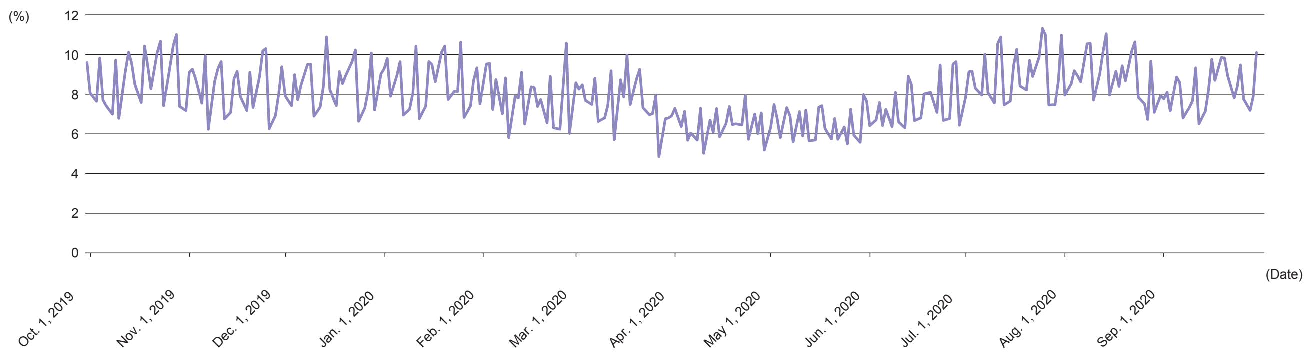

Figure 5: IPv6 as a Proportion of Total Traffic

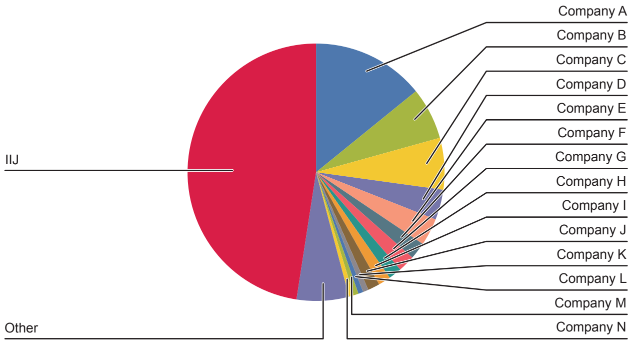

Figure 6: Top Annual Average IPv6 Traffic Source Organizations (BGP AS Number)

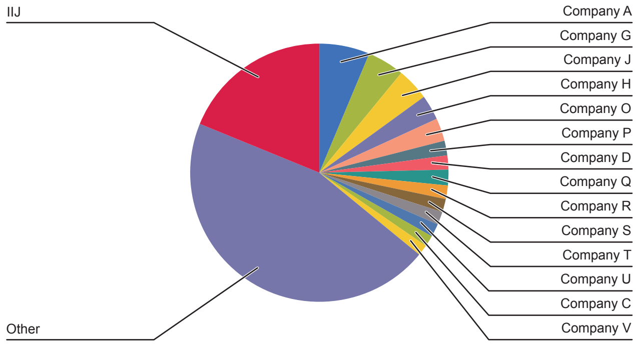

Figure 7: Top Annual Average IPv4 Traffic Source Organizations (BGP AS Number)

Only the top 5 are readily discernible on the graphs. Traffic is scarcer for No. 6 on down and does not show up readily on the graphs.

In the IPv4 space, protocol usage does not appear to have changed much, but nighttime peak traffic has fallen slightly, and daytime traffic appears to have increased overall. With the increase in people remoting into work from home during the day, it looks like daytime traffic has increased or that the data transfer peak has shifted to an earlier time. The shape of the weekday (peaks 1–5) and the weekend (peaks 6 and 7) sections of the graph are very similar, so the differences between weekday and weekend traffic trends appear to have diminished.

Summary

In this issue, we examined IPv6 traffic volume, source ASNs, and protocols used. With the impact of COVID-19, we observed different trends than in the past. IPv6 traffic as a percentage of total usage ultimately did not change much vs. last time, but the impact of people staying at home in the middle of the year was apparent. We observed new trends likely due to remote working in the IPv6 protocol usage figures, the IPv4 traffic peaks, and so on.

The COVID-19 pandemic is yet to show signs of winding down fully. As we move through what are being called the with-COVID and post-COVID eras, we will continue to monitor changes in IPv6 and Internet usage patterns.

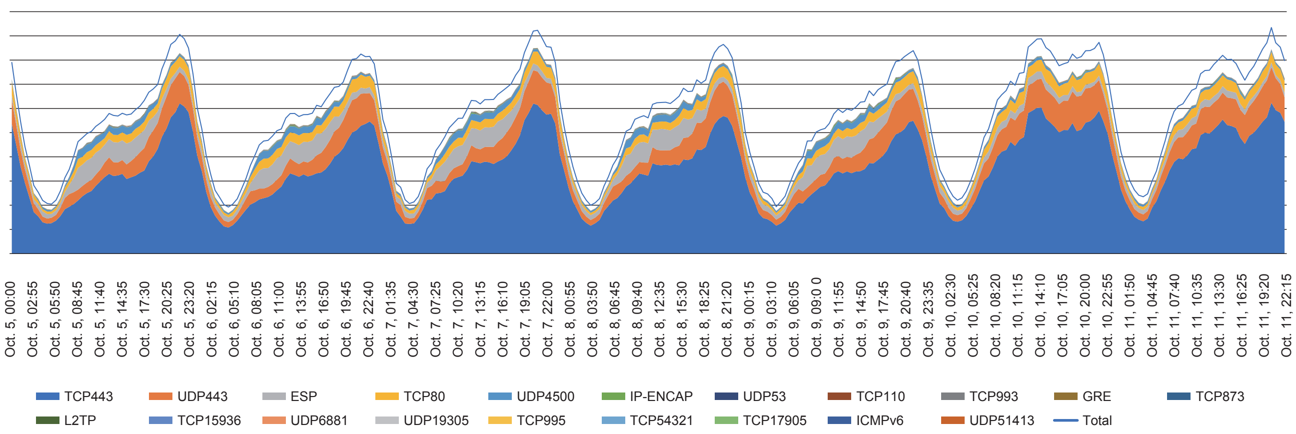

Figure 8: Breakdown of IPv6 Traffic by Protocol Number (Next Header) and Source Port Number

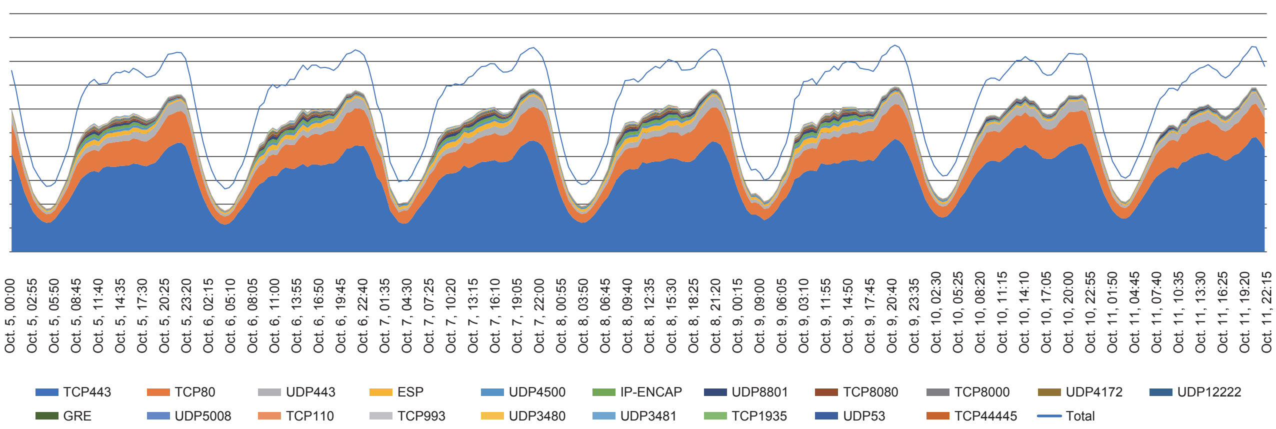

Figure 9: Breakdown of IPv4 Traffic by Protocol Number and Source Port Number

Topic 4
Mobile 3G and LTE

The trends in mobile traffic in 2020, as discussed in Vol. 48 (https://www.iij.ad.jp/en/dev/iir/048.html), were different from usual, with traffic down substantially during periods when people were staying home amid the COVID-19 situation. The term 5G has also gone mainstream in the mobile space, with MNOs launching 5G services, and on October 30 IIJ also began offering a service that supports au 5G as part of the IIJ mobile services that it provides to business customers. While new standards frameworks are gaining traction, old standards are heading for end of life. In October 2019, NTT Docomo announced that its 3G FOMA service will close down at end-March 2026.

In this edition, we look at 3G traffic on IIJ's mobile services based on observations for the period from October 1, 2019 to September 30, 2020.

Figure 10 shows 3G as a proportion of total traffic.

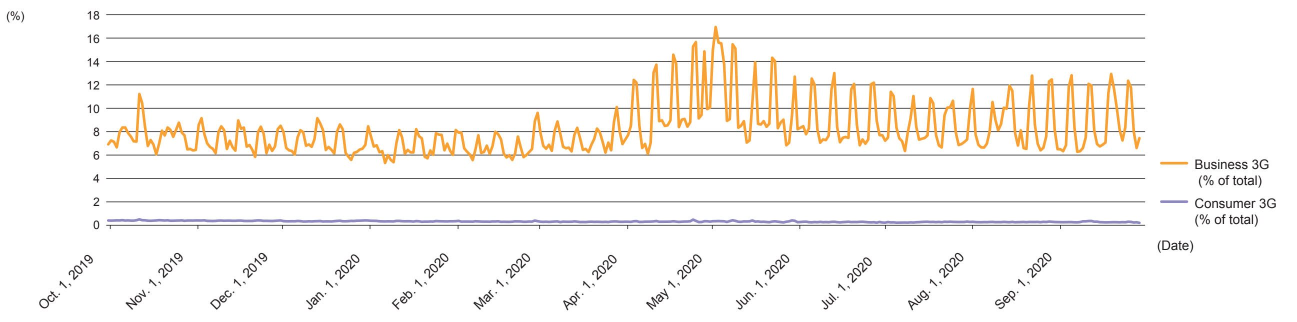

Figure 10: 3G as a Proportion of Total Traffic

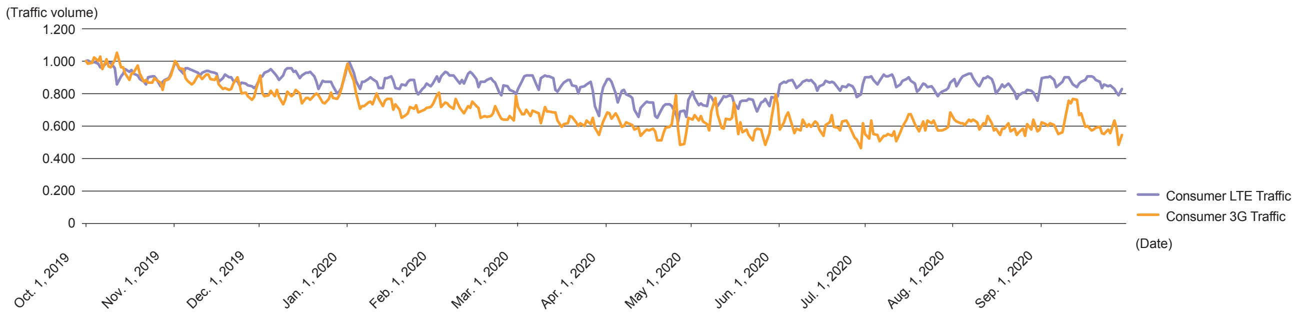

Figure 11: Consumer Service Traffic Volume

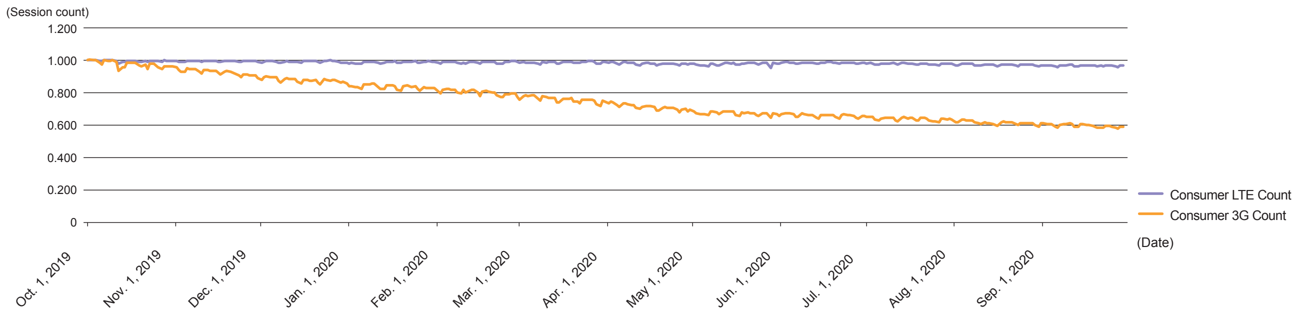

Figure 12: Consumer Service Session Count

On consumer services, 3G is barely used, accounting for under 0.5% of traffic. On business services, meanwhile, it averages around 8% of traffic, indicating that 3G remains well embedded in the corporate space.

Next, we graph trends in consumer service traffic volume (Figure 11) and session count (Figure 12) indexed to October 1, 2019. The figures for traffic volume and session count for 3G on consumer services continue to decline, with both having fallen by around 40% over the past year. One can think of various reasons for this, but given that almost all devices on consumer services are now smartphones, a reasonable explanation is that with the improvements in LTE connectivity at large, connections rarely drop to 3G anymore. We will be keeping tabs on how this unfolds ahead.

Turning to LTE traffic on consumer services, the session count remained largely flat, while traffic volume experienced a lull March through May 2020, when it was down 30%, probably because of people staying home amid the COVID-19 situation.

Next, we look at trends in business service traffic volume (Figure 13) and session count (Figure 14), also indexed to October 1, 2019.

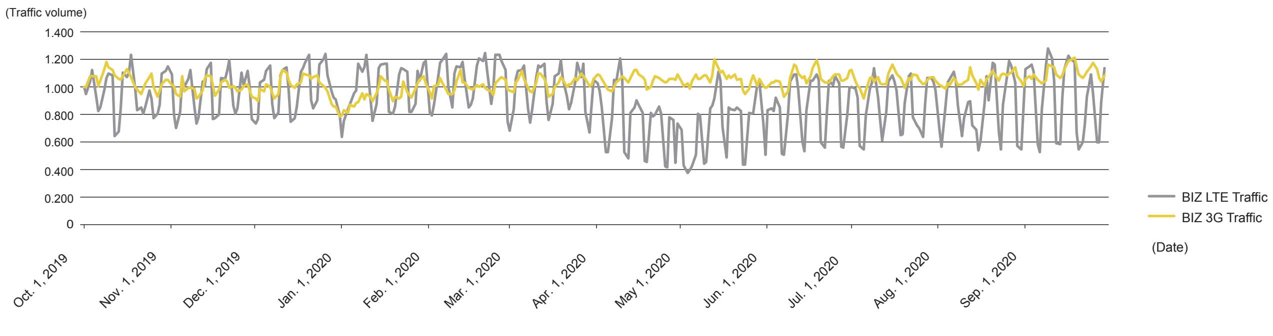

Figure 13: Business Service Traffic Volume

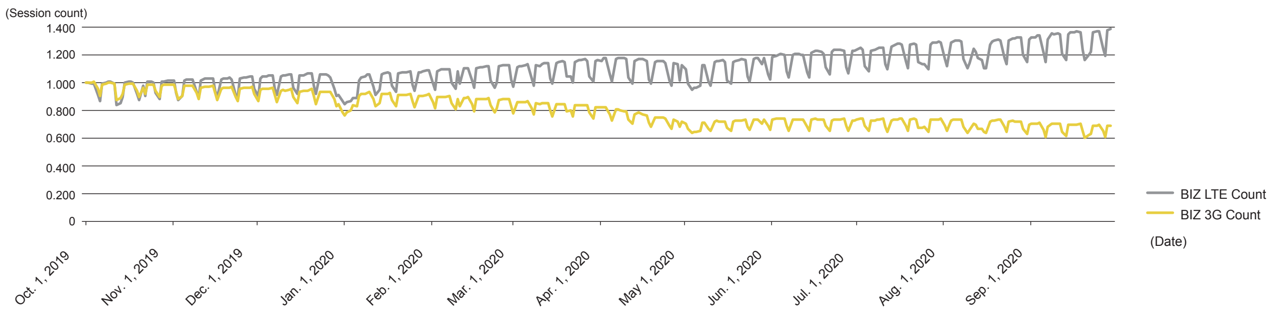

Figure 14: Business Service Session Count

Session count for 3G traffic on business services is in an intermittent downtrend. One likely reason for this is the ongoing migration from 3G to LTE in anticipation of the end of 3G services. Also, the slope of the decline in May 2020 onward is more moderate than from October 2019 up to April 2020, and one possible reason for this is that the speed with which companies are migrating has eased due to the effects of COVID-19. Meanwhile, 3G traffic volume is in a mild uptrend, unaffected by COVID-19. This phenomenon is dependent on the usage of business users, so we will be watching developments closely ahead.

Looking at LTE traffic trends, we see a drop in session counts during the stay-at-home period as with 3G, but the overall trend is of an intermittent rise. Traffic volume troughed in April–May 2020, when people were staying home, and is gradually coming back.

Finally, we look at 5G. As mentioned, IIJ released a business service that uses au 5G on October 30. So while we are not yet able to analyze traffic trends and the like, we investigated the extent to which IIJ users are using 5G-capable devices.

Figure 15 shows the rate of growth in devices likely to be 5G capable (Android devices with 5G in the device name and the iPhone 12 series) relative to October 1, 2019. Over the year before the iPhone 12's release, the figure increased about 40 fold, but in the year after release, it increased around 400 fold (around 10 fold vs. just before release). Traffic trends related to 5G services will also bear close watching ahead.

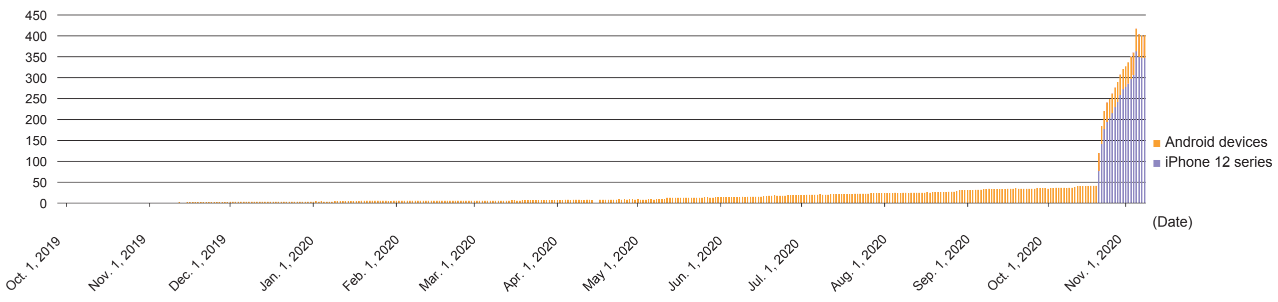

Figure 15: 5G-capable Device Connections

Topic 5

Deploying BGP ROV on the IIJ Backbone

Since November 2020, we have been progressively rolling out BGP ROV (Route Origin Validation), which uses RPKI, on IIJ's Internet backbone.

Note that we also discussed the deployment and workings of RPKI etc. on the engineers blog ahead of the deployment. The number of RPKI ROAs issued, etc., can be looked up via RIPE, NIST, and so on. RPKI itself has been used by RIRs since around 2008, and JPNIC also started taking registrations in 2015. RIPE is a fair way ahead in terms of number of ROAs, but other RIRs' ROA counts have increased greatly in the past few years, so usage appears to be progressing steadily.

Here, we extract information from a specific day in October 2020 before IIJ adopted BGP ROV. Table 5 was created from VRP (Validated ROA Payloads) records from IIJ's ROA cache server for that day in October 2020. You register ROAs with a set of information: a prefix that your organization advertises, an origin AS number that identifies your organization, and a max length, being the maximum prefix length that is acceptable. The number of addresses registered in ROAs as a proportion of BGP routes represents registered ROA prefixes as a proportion of the number of unique BGP route addresses on the specific day for IIJ.

Next, Figures 16 and 17 show the distributions of prefix length and max length in ROA registrations.

Table 5: VRP Data from IIJ's ROA Cache

	IPv4	IPv6	Total count
Unique prefixes	144,785	25,085	169,870
Unique ASNs	16,479	8,769	17,670
Unique prefixes+ASNs	158,099	27,024	185,123
AS0 prefixes registered	184	100	284
No. of ROA-registered addresses as % of BGP routes	27.9%	32.8%	-

Trust Anchors: RIPE NCC, ARIN, APNIC, AfriNIC, LACNIC

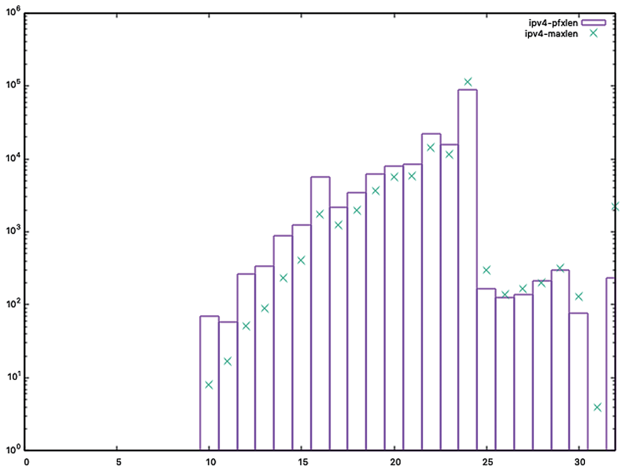

Figure 16: Distribution of Registered Prefix Lengths and Max Lengths (IPv4)

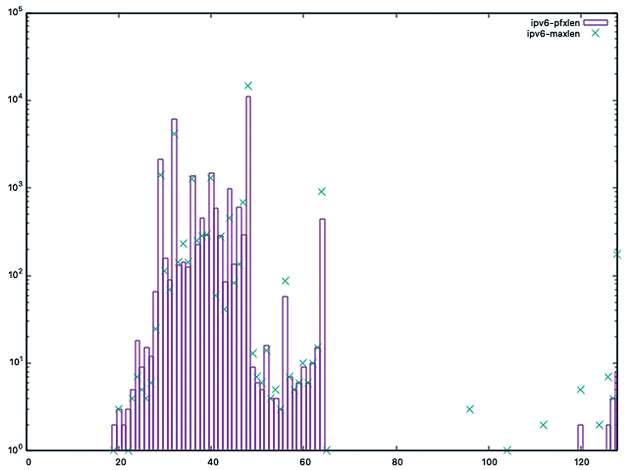

Figure 17: Distribution of Registered Prefix Lengths and Max Lengths (IPv6)

Prefix length is on the horizontal axis and registration count is on the vertical axis. The bars indicate prefix length in the ROA registrations, and the crosses indicate max length, which is the maximum acceptable prefix length. The registered prefix length and max length have the same values in 81.6% of cases for IPv4 and 78.7% of cases for IPv6. Meanwhile, max prefix tends to be longer for the /24 prefix on IPv4 and for the /48 prefix on IPv6, which looks to be the same sort of trend as when organizations exchange BGP route information on the Internet.

Figures 18 and 19 show the extent to which invalid routes are found when validating BGP routes for IIJ's region based on the aforementioned VRP data.

Around 24% of IPv4 and 29% of IPv6 routes are deemed valid, with around 0.32% of IPv4 and 0.49% of IPv6 routes deemed invalid. NotFound, indicating that a prefix matching the ROA was not found, accounted for 70% overall. As ROA registrations increase ahead, the NotFound slice of the pie can be expected to shrink. Even for invalid routes,

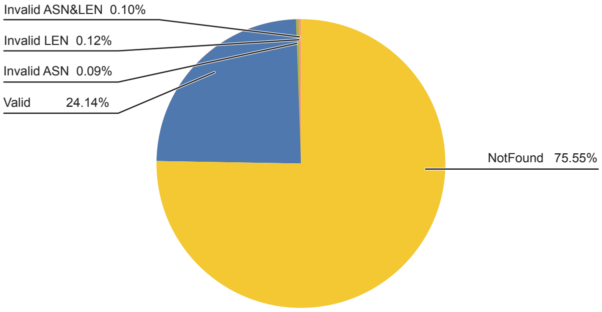

Figure 18: BGP Route Validation Results for a Specific Router (IPv4)

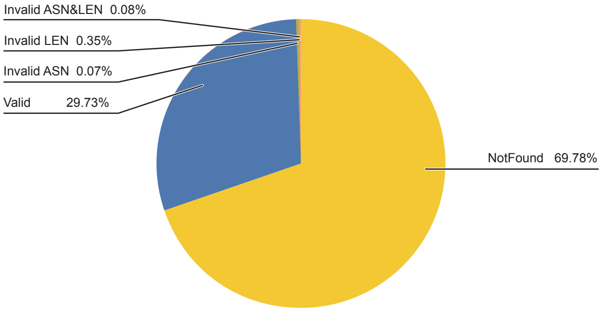

Figure 19: BGP Route Validation Results for a Specific Router (IPv6)

a larger IP space is determined to be valid or NotFound is some cases, so reachability is not necessarily lost if a route is deemed invalid, but there are also routes for which this is not the case, and these account for around 0.028% of IPv4 and 0.02% of IPv6 routes.

The Internet is constantly changing, so we are constantly exposed to the threat of a range of problems with routes due to misconfigurations and malfunctions. Determining whether these sorts of routes are valid or not can be problematic, so it has so far been very difficult to prevent problems from occurring. But with these recent initiatives along with the uptake of RPKI, the possibility of preventing some threats ahead of time is rising.

IIJ will continue to pursue initiatives geared toward providing resilient infrastructure to support a pleasant and convenient Internet experience for all.

1.BGP Routes
Tomohiko Kurahashi
Technology Development Section, Operation Technology Department, Infrastructure Engineering Division, IIJ

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

3.IPv6
Taisuke Sasaki
Deputy General Manager, Network Technology Department, Infrastructure Engineering Division, IIJ

4.Mobile 3G and LTE
Tsuyoshi Saito
Mobile Technology Manager, Network Technology Department, Infrastructure Engineering Division, IIJ

5.Deploying BGP ROV on the IIJ Backbone
Fumiaki Tsutsuji
Network Planning Manager, Network Technology Department, Infrastructure Engineering Division, IIJ