Internet Trends as Seen from IIJ Infrastructure —2021

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 after the deployment of BGP ROV on the IIJ backbone.

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). Incidentally, it was projected at the start of 2021 that APNIC would completely exhaust its IPv4 address pool, but that has not happened over the year that followed.

The annual increase in the number of routes fell short of 40,000 for the first time in 10 years. The increase for prefixes /21 through /24 was also below the previous year's level, and the total number of routes, now over 850,000, may be nearing its peak. The large increase in /8 routes and unique addresses is noticeable. All of the additional /8 routes are advertised by AS8003. Information indicates that the routes advertised by AS8003 (subsequently changed to AS749) are for a special purpose, so excluding the effect of these (765 routes in total), the number of /8 routes was down (-2), and the number of unique addresses was only up slightly (around 2.86 billion).

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. 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
Sep. 2021	16	13	41	101	303	589	1191	2007	13408	8231	13934	25276	41915	50664	106763	91436	497703	853591

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. 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
Sep. 2021	223	3628	705	20650	13050	10233	4170	11545	5204	61024	130432

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).

The total number of routes increased around 1.5-fold to over 130,000, well in excess of the forecast we made last year. The combined total with IPv4 included is now over a million, so keeping network equipment running may have been a challenge for some. The number of unique /64 blocks, meanwhile, only increased by a bit under 5%. The main reason for this seems to be that 58.6% of the total number and 79.1% of the increase were routes for which there were other shorter-prefix routes, meaning that they do not contribute to the increase in the number of unique blocks. We can infer that the IPv6 rollout on end sites has progressed even further, but if unaggregated route advertisements come to dominate the increase in the number of routes going forward, this would be a little disappointing.

Lastly, let's also look at IPv4/IPv6 full-route Origin AS figures (Table 4). In the past year, an additional 6144 32-bit-only AS numbers were allocated to APNIC and 2048 to ARIN.

Both the decrease in 16-bit Origin Autonomous System Numbers (ASNs) and the increase in 32-bit-only Origin ASNs were around double those in the previous year. 32-bit-only ASNs now account for 46.7% of all Origin ASNs, and we expect this figure to exceed 50% in 2022. IPv6-enabled ASNs, which advertise IPv6 routes, also rose substantially to account for 34.5% of the total. Within this, there was a notable increase in 32-bit-only ASNs only advertising IPv6 routes, with 90% of those being APNIC-region ASNs. In 2022, we will be watching to see if this increase in ASNs is a temporary phenomenon amid the IPv6 rollout or if it is an inevitable trend resulting from the recent surge in IPv4 address prices that will continue ahead.

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. 2012	2,588,775,936	41,097,754,610
Sep. 2013	2,638,256,384	20,653,282,947
Sep. 2014	2,705,751,040	62,266,023,358
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

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. 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%)
Sep. 2021	11465	29219	302	40986	(28.7%)	9514	21108	5242	35864	(41.1%)

Topic 2

DNS Query Analysis

IIJ provides a full 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 resolver observational data obtained on October 6, 2021.

The full resolver provides a name resolution function that serves user queries. Specifically, to resolve a name, it 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 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 without relying on OS settings.

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 try to 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 0.12 queries/sec per source IP address at around 4:20 a.m., and a peak of about 0.30 queries/sec per source IP address at around 9:00 p.m. These are 0.06pt increases vs. 2020 in both cases. Comparing the data with the previous year's, the trends are not all that different at times during the day, which is when traffic is higher, but the nighttime trend looks to have shifted, with the number of queries being around 1.8-fold higher. This growth is observed across almost all times of the night, so it may be that automated mechanisms of some kind (e.g., device control, checks of device active status, scheduled tasks) are pushing the figures up.

Broken down by protocol (IPv4 and IPv6), IPv4 queries per IP address rose vs. the previous year. As noted above, automated mechanisms of some kind may be pushing the number of queries up. Changes in implementation and so forth may also be factors. Also, the figures on query source IPs show there to be more IPv6 query source IPs than IPv4 ones during the day, while IPv4 query source IPs are more numerous than IPv6 late at night. The counts are roughly the same around 6:50 a.m. and 10:10 p.m. The increase in IPv4-based queries also affected the overall trend in the number of queries. In a departure from the trend up until 2020, IPv4-based queries accounted for around 59% of the total, while IPv6 accounted for around 41%.

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 20, 14, and 10 seconds before every hour mark, but the increase at the 20-second mark was not very discernible in 2021. Similar to 2019, the 2021 figures show increases 14 and 10 seconds before every hour. At the hour mark, query volume rises sharply and then tapers 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. 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, A records that query the IPv4 address corresponding to the host name and AAAA records that query IPv6 addresses account for around 80% 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 64% are A record queries and 21% AAAA record queries (Figure 1). With IPv6-based queries, meanwhile, AAAA record queries account for a higher share of the total, with around 44% being A record and 36% being AAAA record queries (Figure 2). Compared with the previous year, we observe drops in A record queries of 15 percentage points for IPv4 and 7 percentage points for IPv6. HTTPS-type records, which we started to see in 2020, accounted for some 11% of IPv4 and 18% of IPv6 queries, marking an increase of 9 percentage points for IPv4 and 12 percentage points for IPv6. The trend in HTTPS record queries appears to be correlated with AAAA records, and HTTPS record queries tend to come in at roughly half the volume of AAAA record queries across all times of the day.

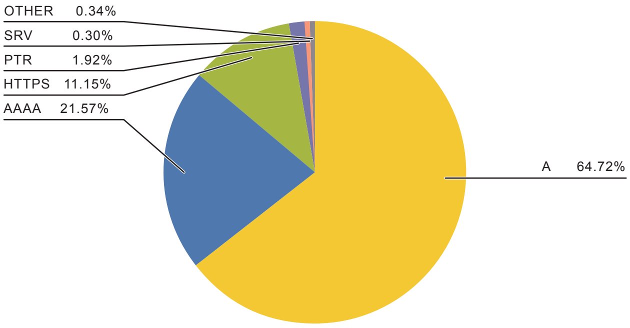

Figure 1: IPv4-based Queries from Clients

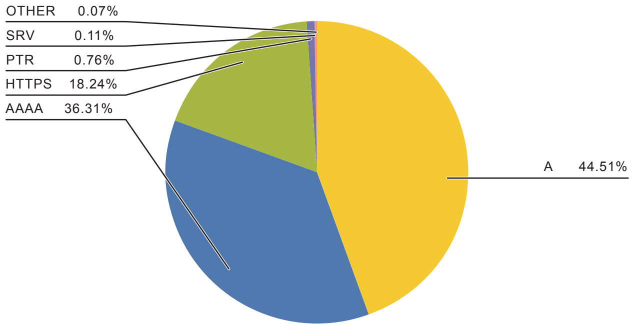

Figure 2: IPv6-based Queries from Clients

Topic 3

IPv6

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

■ Traffic

Figure 3 shows traffic measured using IIJ backbone routers at core POPs (points of presence—3 in Tokyo, 2 in Osaka, 2 in Nagoya). For reasons to do with our system, this edition looks at data for the nine months from the start of 2021 to September 30, 2021, rather than a year's worth of data.

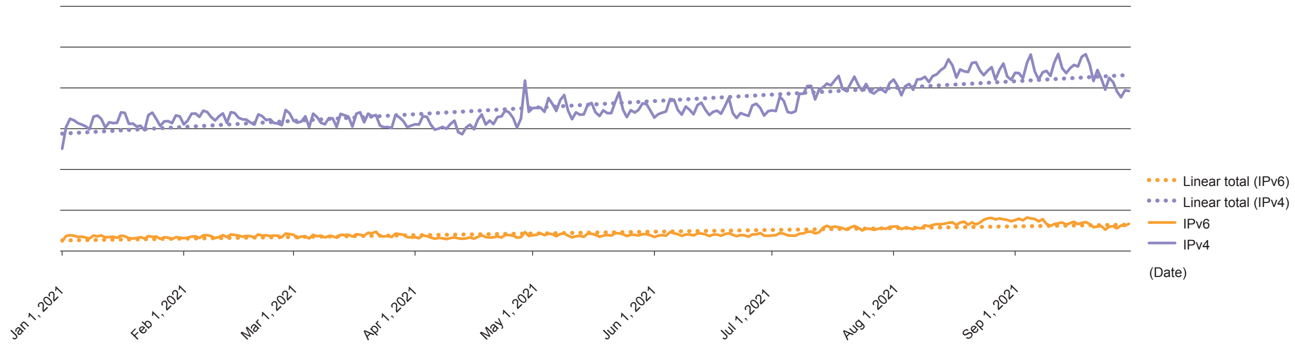

Figure 3: Traffic Measured on Backbone Routes at IIJ's Core POPs

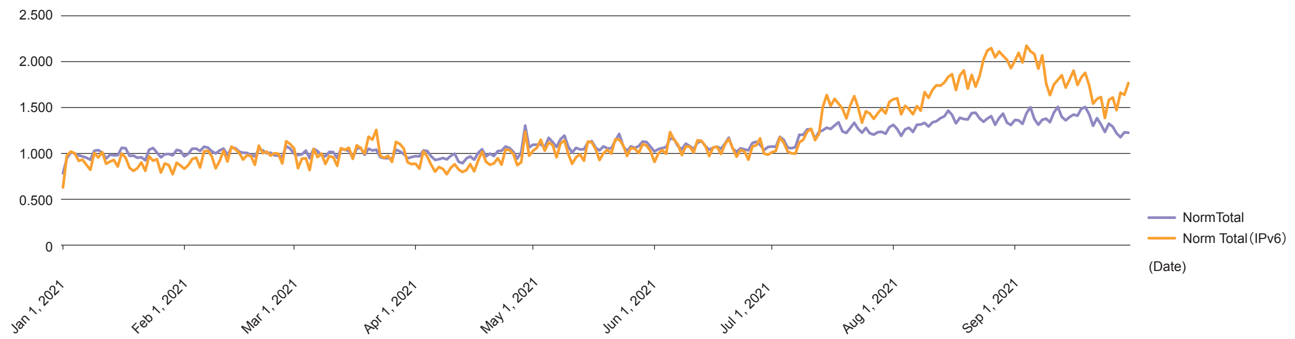

Figure 4: Traffic Indexed to 1 as of January 4

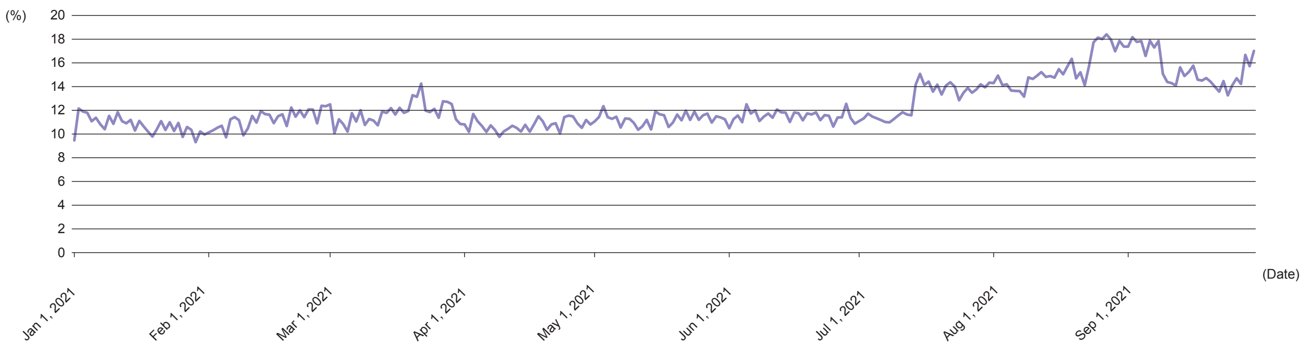

Figure 5: IPv6 as a Proportion of Total Traffic

Possibly on a rebound from 2020's muted growth, both IPv6 and IPv4 traffic volume rose continuously throughout 2021, with particularly strong increases in the latter half of the year. As the normalized series in Figure 4 (indexed to 1 as of January 4, the first business day of the year) show, IPv6 traffic was up 1.7x and IPv4 up 1.2x vs. the start of the year. The graph also shows a lump spanning mid-August through early September. During this period, IPv6 briefly rose to 2.2x and IPv4 to 1.5x vs. the start of the year.

Figure 5 shows IPv6 as a proportion of total traffic. IPv6 traffic continues to increase year after year, but it remains far below IPv4 in absolute terms. But considering that it was essentially crawling along the bottom of the graph back when we started compiling this report in 2017, we can say it has grown considerably over the subsequent four years. With this being the fifth edition of this report, Table 5 recaps the IPv6/IPv4 ratios going back to the first edition.

We observed a slowing of growth in 2020 likely attributable to COVID-19, but the observations this time around confirm that use of IPv6 is rising steadily by the year.

■ 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 January 1, 2021 through September 30, 2021.

Company A retains the top IPv6 spot, with its share of traffic down 3 percentage points vs. 2020 to 11%. The ranking saw a big reshuffling from the No. 2 spot down this time around. Company B in second place with 8% is a major Japanese content company, Company C in third place with 4% is a major US CDN operator, and Company D in fourth place with a 3% share is a major US digital devices maker.

Company B in second place saw traffic rise sharply from around mid-July, rocketing into the No. 2 spot in the space of about two months. One imagines that it embarked on a fairly large campaign to drive the use of IPv6. So with this major Japanese content company finally taking real steps to promote IPv6, similar moves by other companies will no doubt bear close watching ahead.

Table 5: IPv6 as a Proportion of Total Traffic

	IIR Vol. 37, 2018	IIR Vol. 41, 2019	IIR Vol. 45, 2020	IIR Vol. 490, 2021	IIR Vol. 53, 2022
IPv6 ratio	4%	6%	10%	10%	16%

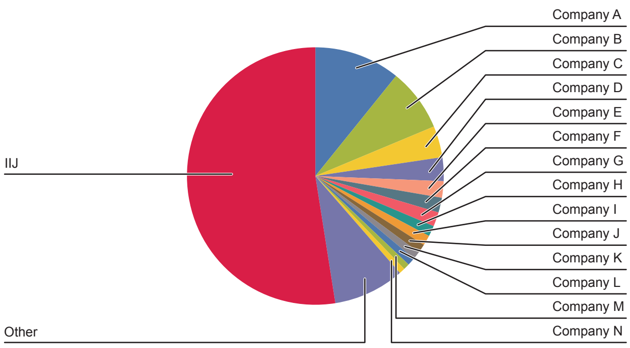

Figure 6: Annual Average IPv6 Traffic by Source Organization (BGP AS Number)

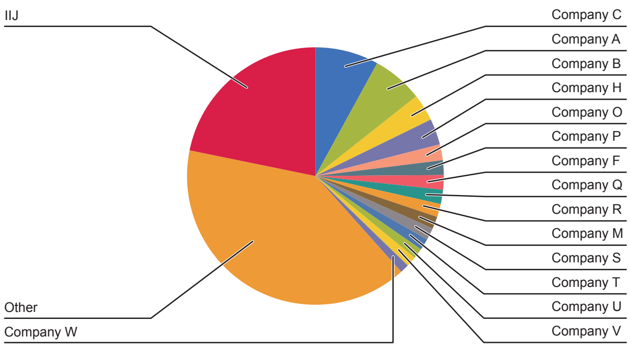

Figure 7: Annual Average IPv4 Traffic by Source Organization (BGP AS Number)

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 of Monday, October 4 – Sunday, October 10, 2021).

The composition of protocols used is largely the same as 2020 for both IPv6 and IPv4. Not much has changed, but TCP 80 continues to decline, and it looks like the transition from HTTP to HTTPS or QUIC is progressing. QUIC was officially enshrined in RFC 9000 in May 2021, so use of the protocol will no doubt grow even more ahead.

We cannot show absolute traffic levels, but IPv6 essentially doubled vs. 2020, with particularly high growth from evening through into night. As mentioned in the discussion on source ASNs, major content company and CDN operator support for IPv6 is progressing, so it's reasonable to think that individual use (games and entertainment) is growing.

Summary

Rebounding from the 2020 doldrums, both IPv6 and IPv4 traffic grew substantially in 2021. IPv6 is also growing steadily as a proportion of total traffic, and breaking through the 20% barrier is certainly within the realm of possibility for 2022. A major Japanese content company appears to have started using IPv6 in earnest, and the level of IPv6 support among CDN operator traffic also seems to be rising.

As a Japanese ISP, we will continue to keep an eye on industry trends in the hopes of seeing a second and then a third major Japanese content company start supporting IPv6 in earnest.

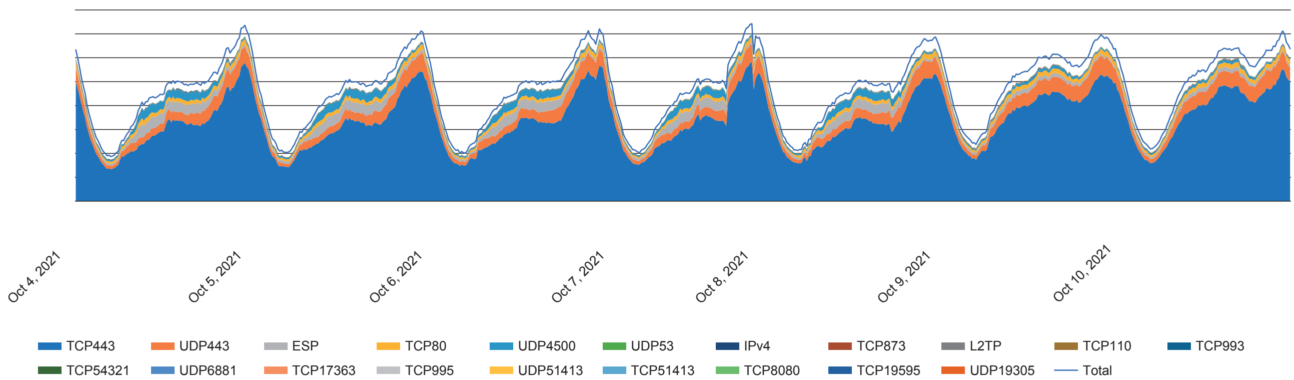

Figure 8: Breakdown of IPv6 Traffic by Source Port Number

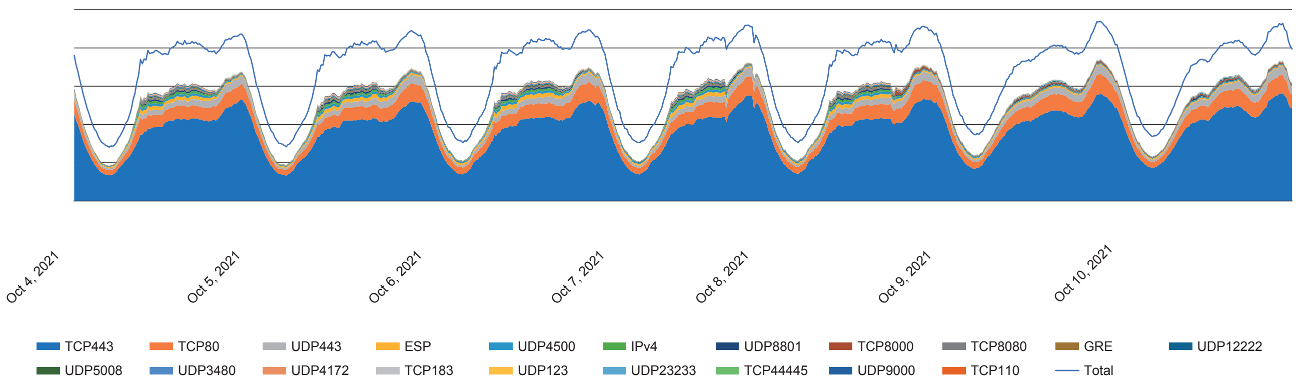

Figure 9: Breakdown of IPv4 Traffic by Source Port Number

Topic 4

State of the Mobile Industry and Traffic Trends

The Suga administration's call for cuts to mobile phone charges has driven the mobile communications industry to review its service offerings in various ways over the past year. IIJ itself released IIJmio Mobile Service GigaPlans ("GigaPlans") on April 1, 2021. As a service provider, we pursue a range of strategies oriented around attracting as many subscribers as we can, but in this section I discuss how we responded in terms of equipment.

From an equipment perspective, our goal was to hold down costs as much as possible while maintaining consistent quality. Two key considerations when it comes to quality are whether and by how much user numbers will rise or fall and how soon we need to react to that forecast. Normally, we also need to take traffic trends (e.g., what time does the peak come each day?) into account, but here we are talking only about IIJmio subscribers, so traffic trends don't enter into it.

User volume forecasts are provided by the department responsible for setting up the service. Usually, based on the user volume forecasts, we look at what bandwidth we need on the interconnection with the MNO and take steps to deal with the speed with which users are set to increase or decrease. To get more precise predictions when releasing GigaPlans, however, we monitored the number of GigaPlans user applications closely to determine a figure for the speed of user growth. As a result, we experienced no major quality issues when GigaPlans was released.

Once GigaPlans was released, our focus turned to building the sort of environment that would not end up being a bottleneck as the services are expanded and enhanced. What we did was review what the PGWs (packet data network gateways), which interconnect with the MNO, were accommodating. IIJ itself has several PGWs running and providing redundancy. We also had plans to add additional PGWs even before plans for the GigaPlans service appeared, so preparations were already underway, but the release of GigaPlans prompted us to work quickly to accommodate IIJmio.

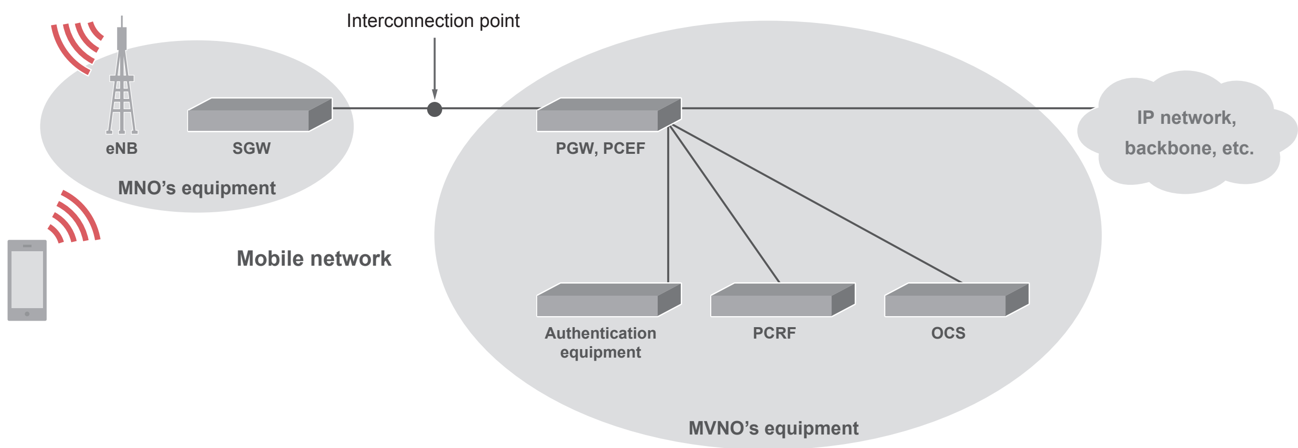

Figure 10: Connection Sequence from Mobile Phone to MVNO Network

Around June 2021 (near the middle of Figure 11), we added a PGW to accommodate IIJmio. As a result, traffic volume increased about 1.4x. We believe this was not simply due to the addition of a PGW but also due to the following factors.

- The addition of accommodating PGWs increased capacity to handle IIJmio processing
- The number of users increased due to the GigaPlans release
- The data limits set on GigaPlans resulted in an increase in traffic per user

That said, we were surprised to see traffic rise so starkly.

We also added a new PGW to accommodate IIJmio around October 2021 (near the right end of the graph). This expansion was equivalent in scale to the one in June, and traffic rose about 1.3 fold vs. immediately before the addition. Here, we believe this reflects an increase in the PGWs' service accommodation efficiency rather than the impact of the rise in user volume immediately following the GigaPlans release.

With mobile equipment, the extent to which you can keep costs down while still providing consistent quality of service is key. Looking ahead, we will consider a range of approaches as needed as we work to enhance mobile service quality.

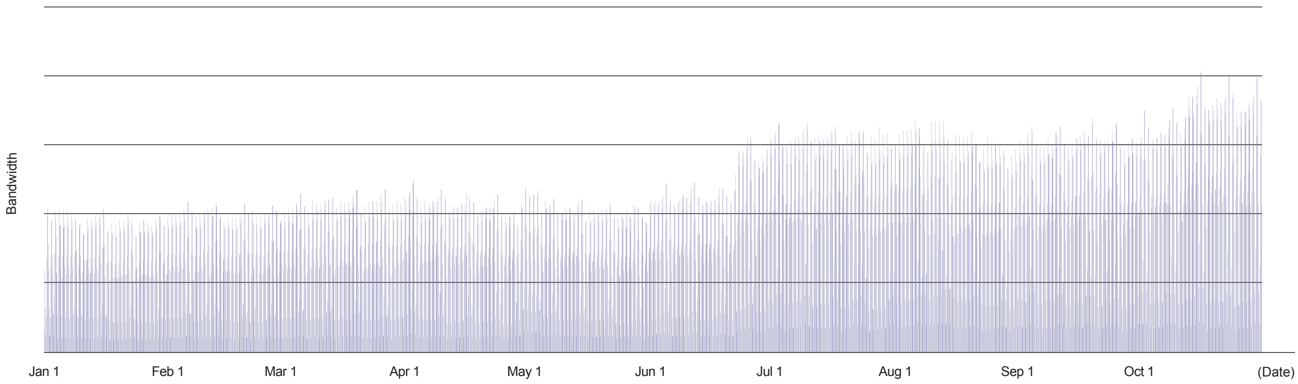

Figure 11:IIJmio Internet-bound Traffic Volume (Jan–Oct 2021)

Topic 5

IIJ Backbone

In the previous edition, we took a look at the situation before IIJ fully deployed ROV using RPKI. We subsequently completed the ROV rollout on IIJ's external connection points as scheduled in December 2021, so roughly a year has now passed.

Comparing what we observed via IIJ's ROA cache server in October 2020 and October 2021, the number of unique prefixes registered was up 1.5x for IPv4 and 1.75x for IPv6. The number of unique Origin ASNs registered was also up 1.49x for IPv4 and 1.4x for IPv6. So we are seeing a steady rise. Comparing the number of unique ASNs registered in ROAs as a proportion of Origin ASNs observable on IIJ's BGP routes, IPv4 was up around 11% and IPv6 up around 1.8% vs. 2020. Meanwhile, looking at the AS list managed by JPNIC (https://www.nic.ad.jp/ja/ip/as-numbers.txt), we see that although the number of ASNs registered as ROAs as a proportion of the JPNIC-managed ASNs that can be seen on IIJ's BGP routes did increase vs. 2020, it still only sits at around 15.5%. IIJ has rolled out ROV, and it is being progressively deployed across the globe as well, so operators that have Internet addresses and other resources can see the effect simply by registering ROAs, and we thus look forward to seeing more and more registrations in Japan going forward as well.

The Tokyo Olympics took place in 2021, so I would also like to look at IIJ's traffic figures during the event. The Games were postponed from 2020 because of COVID-19, and it was an unusual affair in that most of the events, including the opening ceremony, were held without spectators in attendance. These sorts of huge events are major considerations even for communications carriers not directly involved in them. Naturally, we have to be prepared for all sorts of incidents, but the availability of easy access to rich content in recent times means that we also need to watch out for trends in traffic that differ from what we see normally as well as sudden fluctuations in traffic. There is no clear indication as to how much of our infrastructure we will need to have ready, so we have to make capacity available based on estimates derived from the usual traffic trends. On top of that, with these Games, there was a string of announcements right before the opening date about the decision to hold events without spectators and to cancel public viewings. So we anticipated an increase in the number of people viewing the Games on TV at home, or via live feeds on the Internet, and so forth, so we set about preparing for that. In IIJ's case, this involved increasing the capacity of some of the equipment that we already had at the ready ahead of the Games and coordinating with other (non-IIJ) ISPs.

Here, we present graphs overlaying the traffic data for each of the five weeks from July 11 to August 14, 2021, specifically looking at Tokyo/Osaka traffic as well as traffic on external connections with IIJ within Japan (Figures 12 and 13). The data for the period of the Games (July 21 – August 8) are plotted as solid lines; the remainder of the data appear as dotted lines.

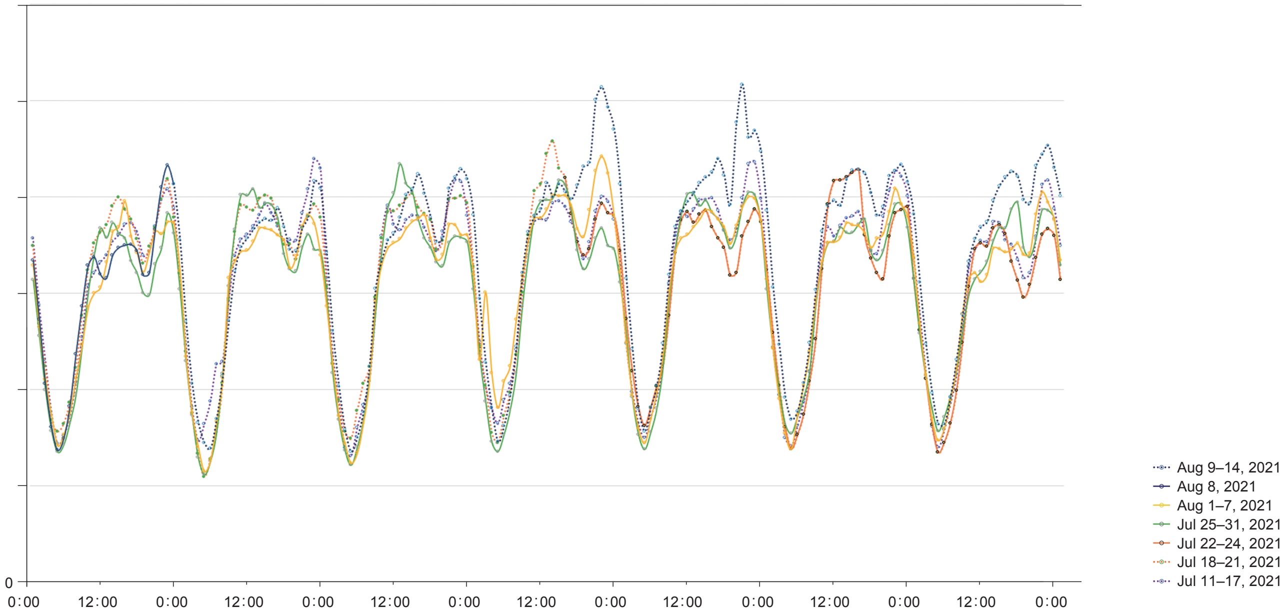

Figure 12: Tokyo/Osaka Traffic

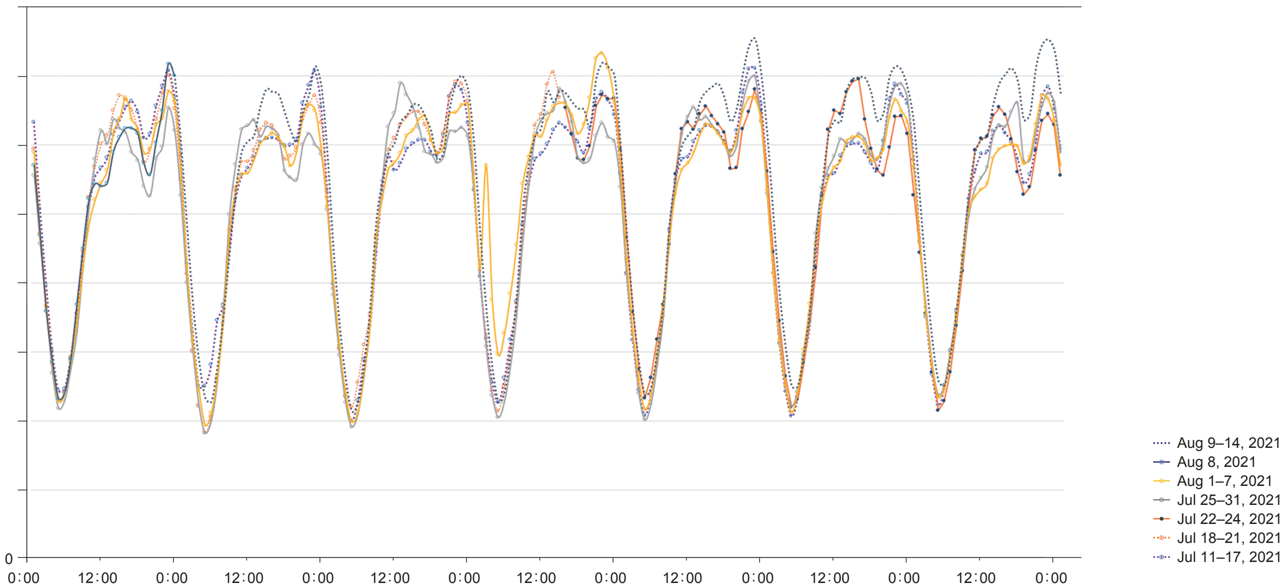

Figure 13: Domestic Interconnection Traffic

As the graphs show, we saw no major increases or decreases in traffic during the Games. The patterns are not all the same, but there was no major deviation from normal traffic levels in the daytime, while traffic during the usually busy 6:00 – 11:00 p.m. period was a few percent below its usual level. Tokyo/Osaka traffic during the 8:00 p.m. timeslot when the opening ceremony took place on July 23 was around 17% lower than during that time on the same day of the week before. It may be that normal Internet usage fell during this timeslot with a few more people watching the opening ceremony on TV.

We made all sorts of preparations heading up to the Olympics, including the ROV rollout mentioned above, but looking back now that it's all over, it was a fairly calm period with no huge changes in traffic from the norm and no major incidents or faults during the Games. We plan to continue expanding our infrastructure to ensure a stable user experience, not just during events like this, but at all times.

1.BGP and 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.State of the Mobile Industry and Traffic Trends
Tsuyoshi Saito
Mobile Technology Manager, Network Technology Department, Infrastructure Engineering Division, IIJ

5.IIJ Backbone
Fumiaki Tsutsuji
Network Planning Manager, Network Technology Department, Infrastructure Engineering Division, IIJ