Internet Trends as Seen from IIJ Infrastructure —2019

To provide Internet services, IIJ operates some of the largest network and server infrastructure in Japan. Here, we examine and discuss current Internet trends based on information obtained through the operation of this infrastructure.

We cover the topics of network routing information, DNS query information, and IPv6 usage, as well as the impact of natural disasters on mobile and FLET'S connection services. We also report on the current state of the backbone network that supports the bulk of IIJ's traffic.

Topic 1

BGP / Number of Routes

We start by looking at IPv4 full-route information advertised by our network to other organizations (Table 1). This time around, we also show the number of unique IPv4 addresses contained in the IPv4 full-route information (Table 2). During the past year, the maximum size of IPv4 addresses allocated by APNIC (and JPNIC) fell to /23 (512 addresses).

The total number of routes, while seeing a slightly smaller increase than in the previous year, now exceeds 760,000. Together, the /22, /23, and /24 prefixes account for 80.1% of all routes. Meanwhile, the number of unique IPv4 addresses, although accounting for less than 1% of the total, fell for the first time in the past nine years. Whether this is a temporary effect, perhaps due to the removal of unauthorized routes using RPKI, or the first sign that the IPv4 Internet is shrinking is something that will bear close watching ahead.

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. 2010	20	10	25	67	198	409	718	1308	11225	5389	9225	18532	23267	23380	30451	29811	170701	324736
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

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

Date	No. of IPv4 addresses
Sep. 2010	2,277,265,152
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

Next we take a look at IPv6 full-route data (Table 3). The total number of routes increased by more than it did in the previous year and now exceeds 70,000. However, route advertisements for blocks that have been split into smaller fragments still constitute the majority, with the top three year-on-year growth rates coming in the /30–/31, /41–/43, and /45–/47 prefix ranges, sizes that are not all that commonly allocated/assigned.

Lastly, let's also take a look at IPv4/IPv6 full-route Origin AS figures (Table 4). Both the decrease in 16-bit Origin Autonomous System Numbers (ASNs) and the increase in 32-bit-only Origin ASNs were the largest seen in the past nine years. And for the first time, IPv6-enabled ASNs, which advertise IPv6 routes, accounted for over a quarter of the total. It was predicted that RIPE NCC's IPv4 address pool would run out before 2020 arrived, so it will be interesting to see what has happened when we next report on the data.

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 October 25, 2019.

The full resolver starts by looking at the IP address of an authoritative name server for the root zone (the highest level zone), and based on the information available from that server, it then goes through other authoritative name serves 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 also implemented on

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. 2010	38	3	10	2023	33	2	9	4	17	436	2575
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

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. 2010	2083	32399	67	34549	(6.2%)	17	478	3	498	(4.0%)
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%)

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.

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 users to automatically configure which full resolver to use for name resolution on their 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 0.05 queries/sec per source IP address at around 4:30 a.m., and a peak of about 0.23 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 in the peak of 0.01 points. Broken down by protocol (IPv4 and IPv6), the trends in query volume are virtually the same (no major differences) in the middle of the night, whereas IPv6 queries per IP address show a tendency to rise when people are active during the daytime and particularly after 8:00 p.m. This suggests that the computing environment needed to allow use of IPv6 in the home is coming into place. 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 60% of the total, up by more than 4 points from 55% 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, which tells us that this 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 an increase in queries 14 seconds before every hour mark, and the 2019 results also show another increase 10 seconds before every hour. The increase in queries that occurs on the hour tapers off gradually, but with the spikes that occurs 14 and 10 seconds before the hour, query volume immediately returns to about where it had been. Hence, because a large number of devices are sending queries in almost perfect sync, it seems like some sort of lightweight, quickly completed tasks are being executed.

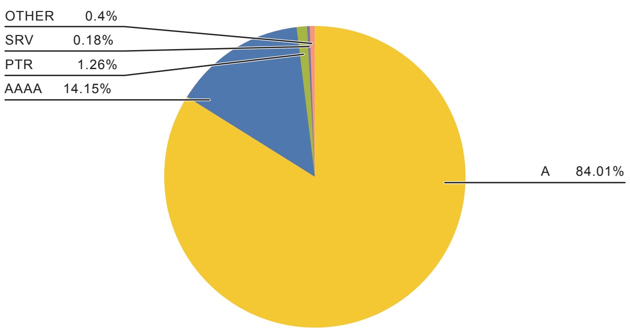

Figure 1: IPv4-based Queries from Clients

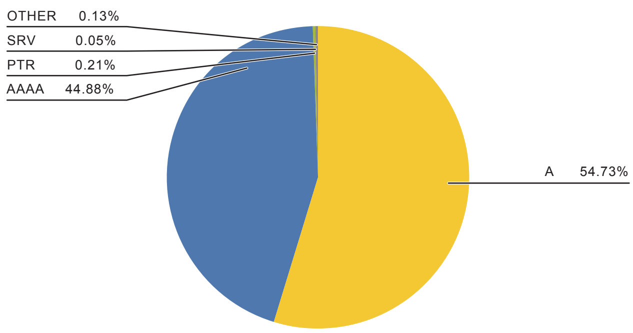

Figure 2: IPv6-based Queries from Clients

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 84% are A record queries and 14% AAAA record queries (Figure 1). With IPv6-based queries, meanwhile, AAAA record queries account for a higher share of the total, with around 54% being A record and 44% being AAAA record queries (Figure 2). Compared with the previous year, the levels are virtually the same for IPv6, while for IPv4, we observe a drop of 3 percentage points or so in A record queries and a rise of 3 points or so in AAAA record queries.

Topic 3

IPv6

In this section, we report on the volume of IPv6 traffic on the IIJ backbone, the sources of that traffic, and the main protocols used. And for a new perspective on IPv6 in the mobile space, we look at the state of IPv6 connections according to differences in device OS (Apple iOS / Android).

■ 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, 2018 to September 30, 2019.

Over the year, IPv4 traffic increased by around 8% while IPv6 traffic rose by around 85%. IPv6 accounts for around 10% of overall traffic (Figure 4), a 4-point increase from around 6% last year. Further, IPv6 traffic hit a peak of

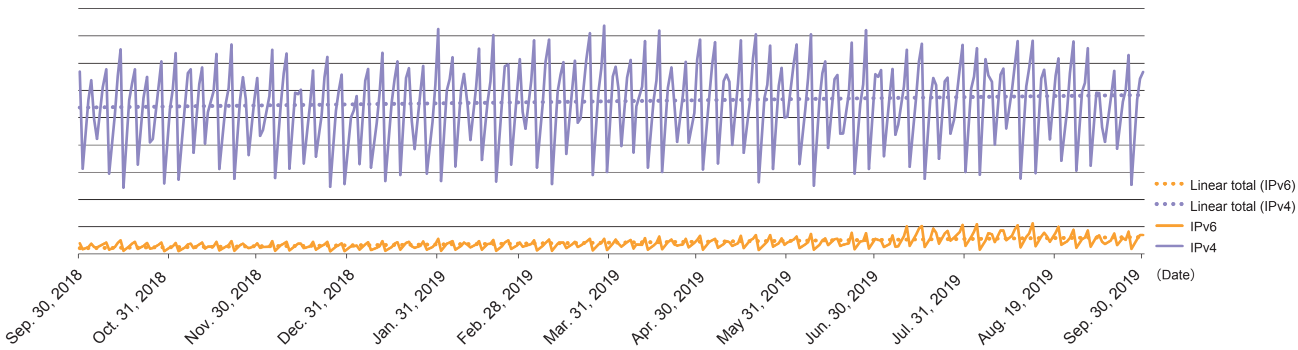

Figure 3: IPv4/IPv6 Traffic Measured via IIJ Backbone Routers at Core Points of Presence (Tokyo, Osaka, and Nagoya)

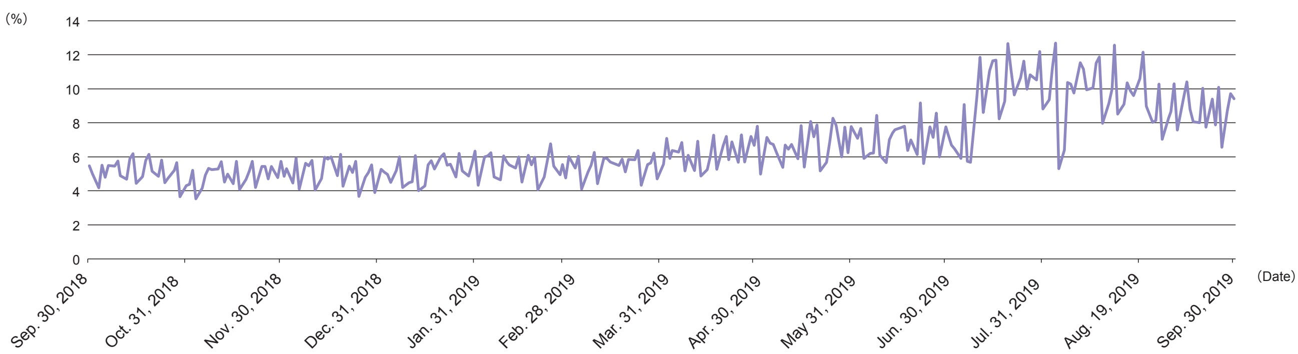

Figure 4: IPv6 Traffic as a Proportion of Total

around 12% of total, and it continues to increase steadily as a proportion of overall traffic.

Figure 5 plots the data for the same period on a log scale. It is evident that IPv6 traffic is growing steadily while growth in IPv4 traffic is slowing.

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 2018 through September 2019.

Company A retains the top spot, but the traffic volume gap between it and No. 2 downward has narrowed further, and its share of the pie is only about 60% of what it was last time. This reflects the increasing use of IPv6 among many operators as well as an increase in IIJ's IPv6 traffic due to IPv6 being enabled on video streaming services that use JOCDN's platform (JOCDN is an IIJ affiliate that provides a video streaming platform).

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 September 30, 2019).

In the IPv6 space, TCP 80 (HTTP) moved from No. 3 last time to No. 2, and UUD 443 (QUIC) came in at No. 3. This mirrors the IPv4 ranking, so we can now say that IPv6

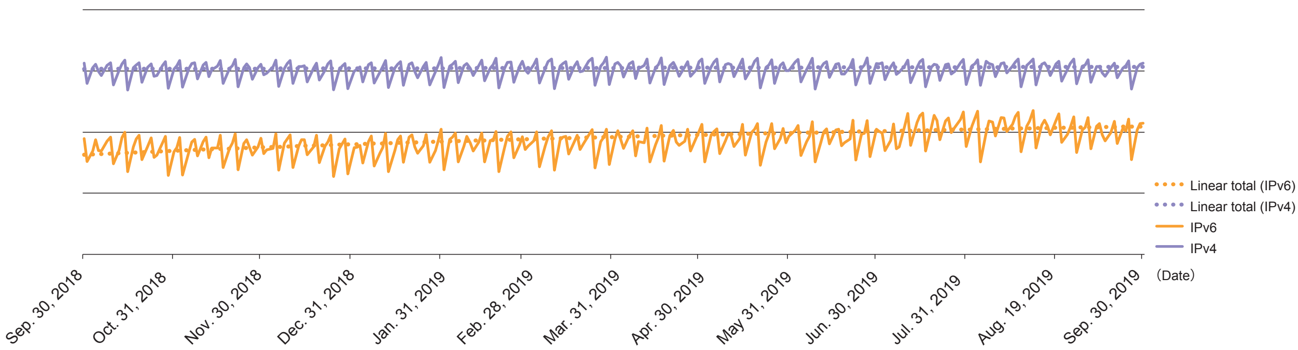

Figure 5: IPv4/IPv6 Traffic Measured via IIJ Backbone Routers at Core Points of Presence (Tokyo, Osaka, and Nagoya)—Log Scale

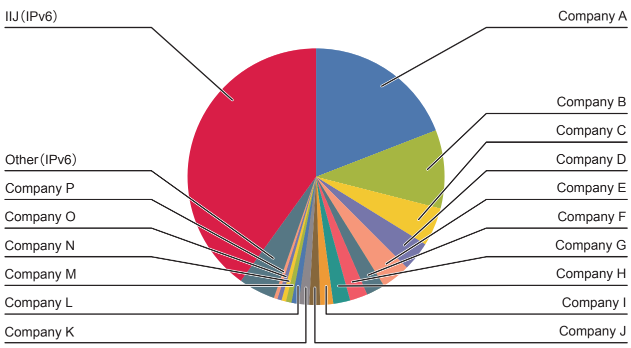

Figure 6: Top Annual Average IPv6 Traffic Source Organizations (BGP AS Number) from October 2018 to September 2019

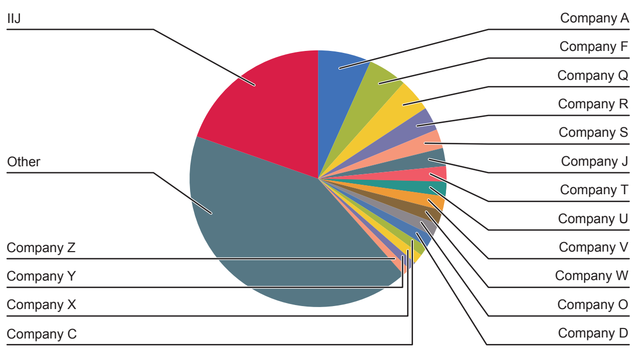

Figure 7: Top Annual Average IPv4 Traffic Source Organizations (BGP AS Number) from October 2018 to September 2019

usage is similar to IPv4 usage or, in other words, that it has become mainstream.

UDP 4500, which was outside the rankings last time, came in at No. 5. NAT is basically not used with IPv6, and it is curious to note that UDP 4500, which is generally for IPSec NAT traversal, ranks toward the top.

■ IPv6 Across Different Device OSs

Last time, we noted that IPv6 is enabled by default in Apple iOS from version 11, and that mobile IPv6 traffic had thus increased.

This time around, we look at mobile device IMEI (International Mobile Equipment Identity) numbers to facilitate an analysis of OS type (iOS or Android) and whether the device is connecting via IPv6 or not (whether an IPv6 address is assigned).

The analysis covers around 1.07 million mobile phones lines on a personal mobile service (IIJmio Mobile Service, number of lines subscribed to as of end-June 2019) that were connected on a particular weekday in October 2019 (MVNE lines and business lines are not included).

First, IPv6 was enabled on 48% and disabled on 52% of connections, a fairly even split, as shown in FIgure 10. While the connections are split fairly evenly, the traffic is about 80% IPv4 and 20% IPv6.

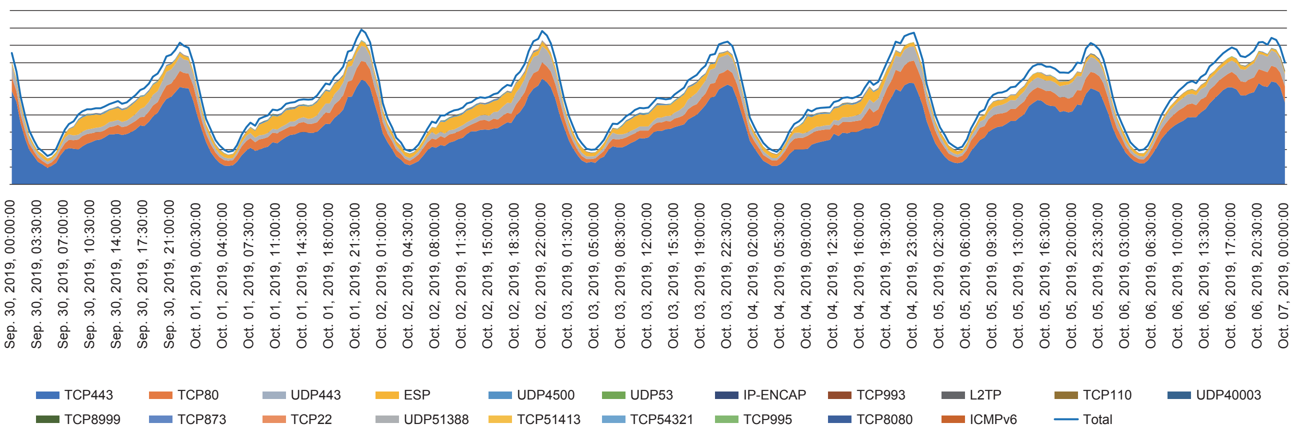

Figure 8: Breakdown of IPv6 Traffic by Protocol Number (Next Header) and Source Port Number

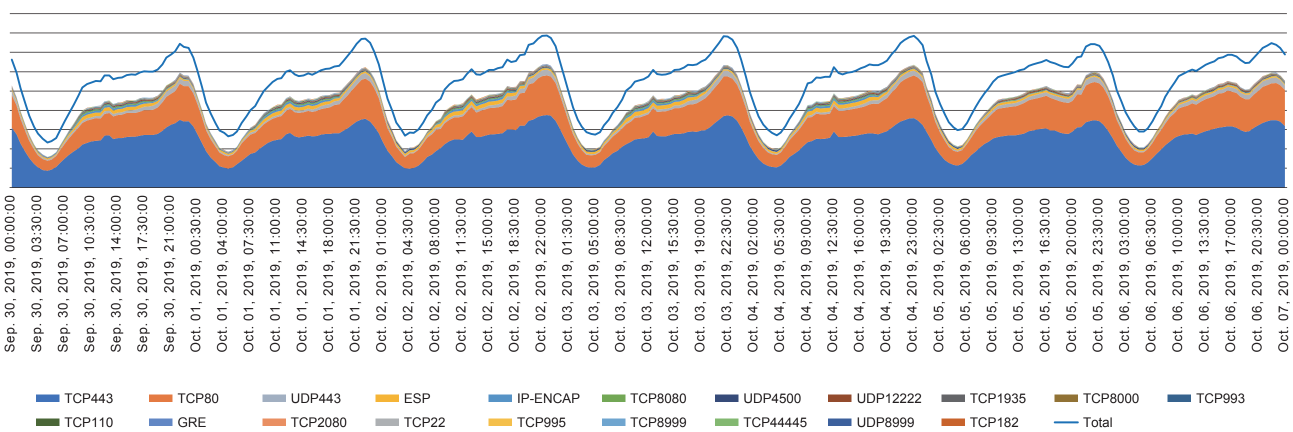

Figure 9: Breakdown of IPv4 Traffic by Protocol Number and Source Port Number

We first look at which OS the IPv6-enabled devices are using. As Figure 11 shows, over 80% are on Apple iOS while 14% are on Android. And next we look at the OS on devices where IPv6 is disabled. As Figure 12 shows, the proportions here are inverted, with Android on 82% and iOS on 8%.

Android has supported IPv6 since the release of Android 5 in 2014, much earlier than Apple enabled IPv6 support with iOS 11, but IPv6 is often disabled by default on Android per device manufacturer and MNO policies. So from an IPv6 viewpoint, a stark difference has arisen between Android and iOS, which is exclusively controlled by Apple.

Nonetheless, IPv6 is disabled on some Apple iOS devices; 9.21% of all iOS devices to be precise. We would guess that these are older devices that do not support iOS 11 or higher. IPv6 is enabled on 14.08% of all Android devices, and it appears that many of the recent SIM-lock-free devices have IPv6 enabled.

Summary

In this issue, we examined IPv6 traffic volume and protocols used, as well as IPv6 connection rates based on mobile device OS. While IPv4 traffic growth is slowing, IPv6 traffic exhibited similar levels of growth to last time, indicating that use of IPv6 continues to advance. Many recently available home Wi-Fi routers offer support for IPv6 IPoE, and there are moves to enable IPv6 on video streaming services as well, so IPv6 traffic looks set to rise even further.

IPv6 is available in the mobile space more than we had imagined, being enabled on almost half of connections. It still only accounts for around 20% of all traffic, so we hope that services continue to provide even more IPv6 support ahead.

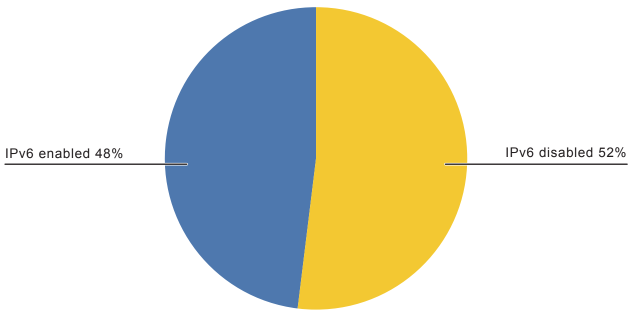

Figure 10: Proportion of Connections with IPv6 Enabled

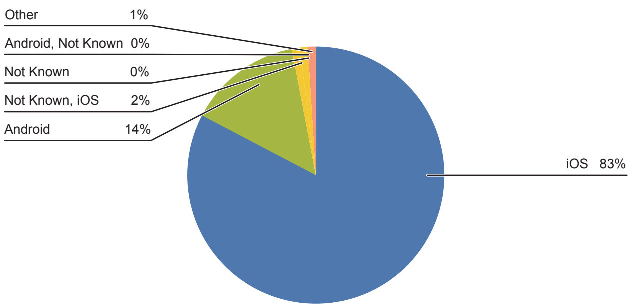

Figure 11: OS Breakdown for IPv6-enabled Devices

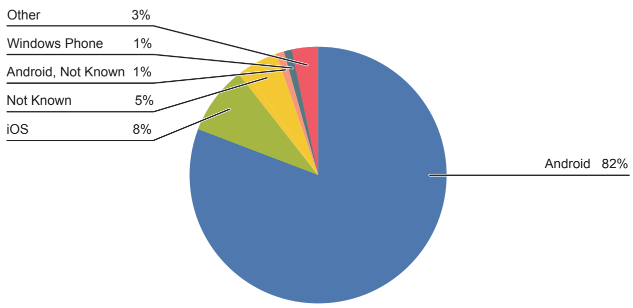

Figure 12: OS Breakdown for IPv6-disabled Devices

Topic 4

Mobile FLET'S and Natural Disasters

Usually, we would report on changes in the trends in FLET'S and mobile traffic based on comparable observations, but virtually the same content was presented in our periodic observation report in Vol. 44 (https://www.iij.ad.jp/en/dev/iir/044.html), so here we look at the effects of natural disasters.

Many natural disasters occurred in 2019. Typhoon Hagibis, in particular, wrought serious damage across many areas. On the IIJ backbone, several lines connecting Tokyo and Osaka experienced lengthy disconnections. A road was washed away in Nagano Prefecture (between Tokyo and Osaka), and buried optical fiber cables were physically damaged. The damage spanned a long section of the route, and the recovery took weeks. Fortunately, IIJ's backbone network is designed to handle faults caused by events such as natural disasters like this. The Tokyo–Osaka leg is distributed among Pacific Ocean, Sea of Japan, and inland routes, so there was no real impact on user communications.

Meanwhile, Typhoon Hagibis had a clear impact on access services linked directly users' use of mobile and FLET'S services. Let's examine this with reference to the data. Note that FLET'S here refers only to PPPoE.

Figures 13 to 24 graph number of connections and traffic volume around October 12, 2019, when Hagibis hit Japan's main island, along with the corresponding data for a week earlier around October 5, 2019 for comparison (upload and download). All of the series have been indexed to a value of 1 at 12 a.m. on the Fridays (October 4 and 11, 2019).

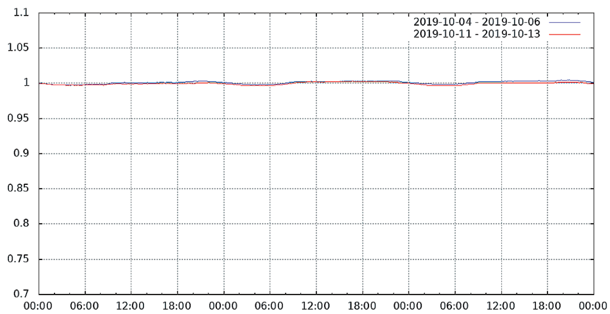

Figure 13: Number of FLET'S Connections in the Osaka Area

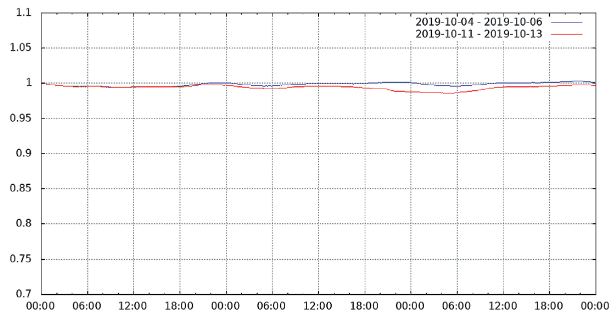

Figure 15: Number of FLET'S Connections in the Tokyo Area

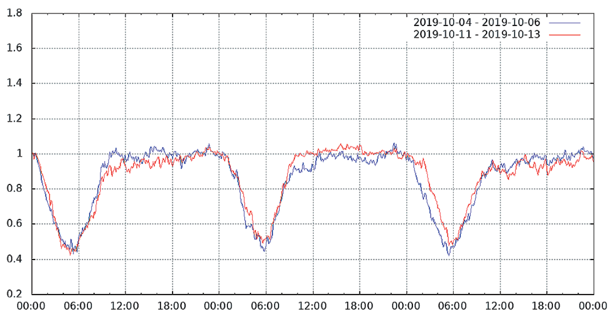

Figure 14: FLET'S Traffic Volume in the Osaka Area

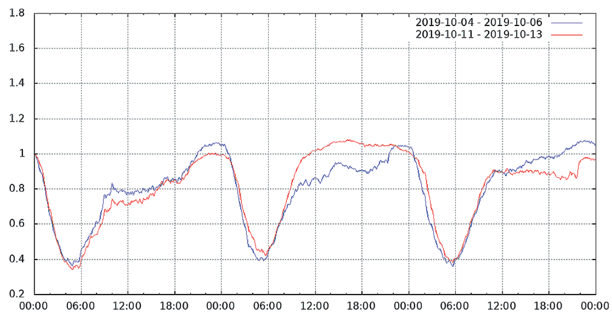

Figure 16: FLET'S Traffic Volume in the Tokyo Area

First, we focus on the number of connections. Currently, most FLET'S users have always-on connections via broadband routers and home gateways, so connection volume remains fairly constant throughout the day. Indeed, connection volume changed only slightly during the period in the Osaka area, which was largely unaffected by Typhoon Hagibis.

However, the number of FLET'S connections in the heavily impacted Chiba, Fukushima, and Miyagi areas declined over October 12 and 13, 2019. It dropped sharply in the FLET'S Chiba area in particular. These areas experienced many power outages, so the drops were likely due to household broadband routers and home gateways going offline when the power dropped out. For mobile, meanwhile, the nature of MVNO equipment precludes the ability to determine device location, so we use nationwide totals. The data show that the number of connections fell over October 12 and 13, 2019. Possible explanations are that MNO equipment

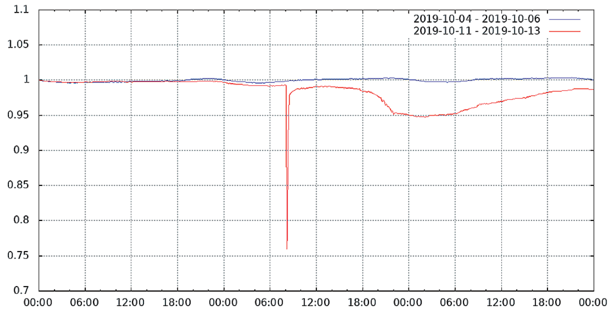

Figure 17: Number of FLET'S Connections in the Chiba Area

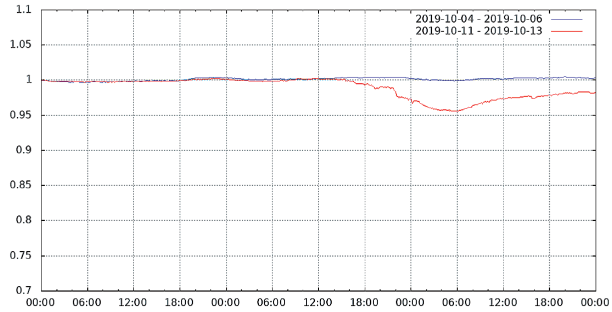

Figure 19: Number of FLET'S Connections in the Fukushima Area

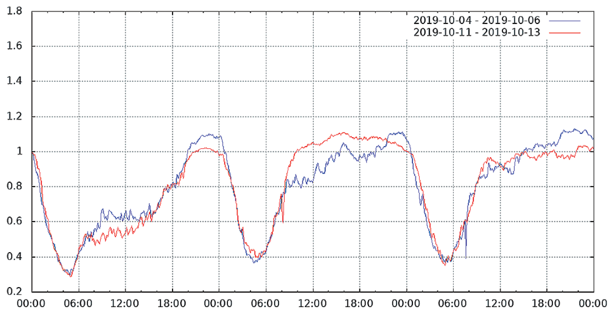

Figure 18: FLET'S Traffic Volume in the Chiba Area

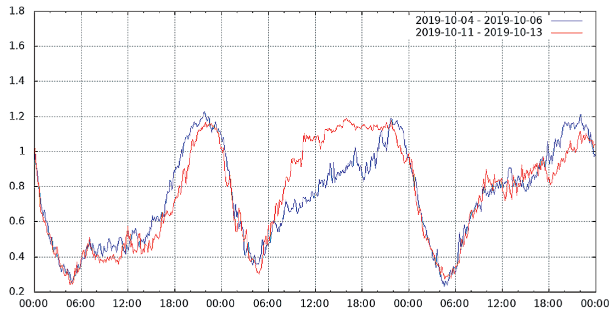

Figure 20: FLET'S Traffic Volume in the Fukushima Area

was also damaged or that people refrained from going out and thus did not use their mobile connection.

Next we look at traffic volume. In Tokyo and areas to the north, FLET'S traffic volume increased during the daytime on October 12, 2019, while mobile traffic fell heavily. It was a weekend, but many people probably stayed indoors to avoid the storm, using their FLET'S broadband at home instead of mobile data.

The Internet is now an important part of our infrastructure alongside electricity, gas, and water, and one that is all the more valuable during natural disasters. IIJ will continue to design its backbone with due consideration to the potential impact of natural disasters and provide Internet infrastructure that users can rest assured will be available at all times.

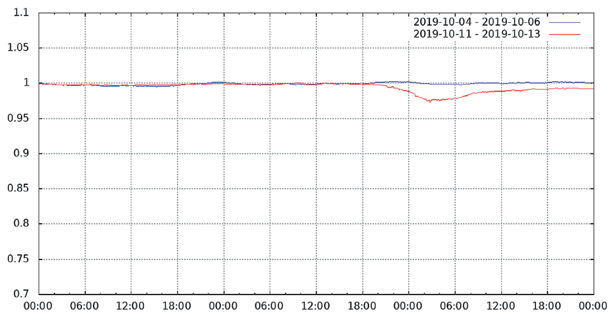

Figure 21: Number of FLET'S Connections in the Miyagi Area

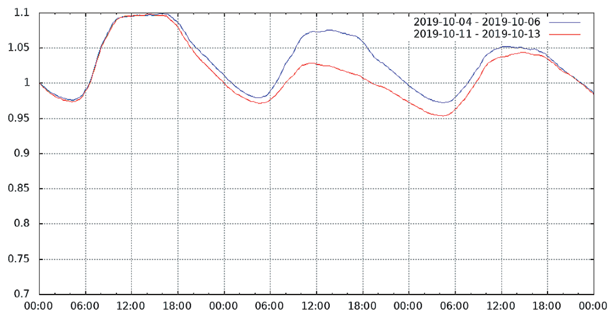

Figure 23: Number of Mobile Connections

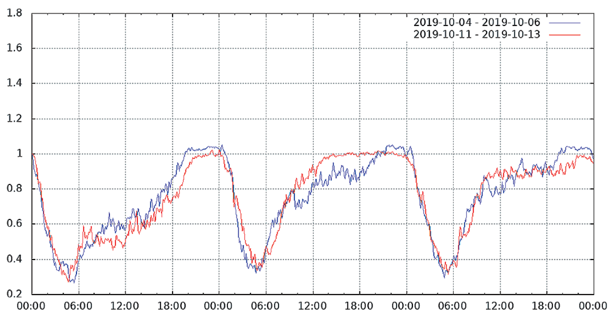

Figure 22: FLET'S Traffic Volume in the Miyagi Area

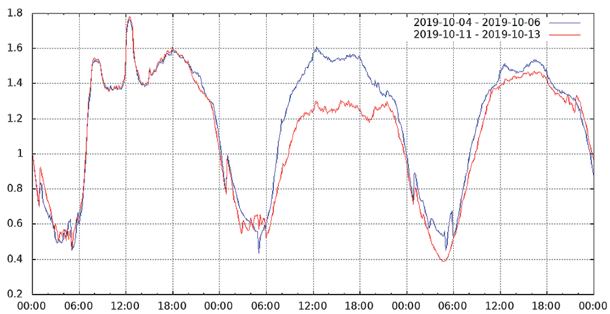

Figure 24: Mobile Traffic Volume

Topic 5

The Transition of the Bandwidth of IIJ Backbone Circuits

In the 2018 edition (Vol. 41) of this report, we noted that total traffic on IIJ's backbone had grown more than tenfold in the past 10 years. Here, we look back on the transition of IIJ's backbone, particularly the bandwidth of the backbone circuits. IIJ began its ISP operations as a Type 2 Telecommunication Carrier; License for the carrier providing telecommunication services using facilities or circuits leased from Type 1 Telecommunication Carriers. So from the outset, we have leased circuits installed by Type 1 Telecommunication Carriers to build a network and provide our services. And this situation is unchanged even after the repeal of the distinction between Type 1 and Type 2 carriers.

We still have the backbone map from 1998. Most of the circuits' bandwidth used at that time were 45Mbps (DS-3), with some being 155Mbps (STM-1). The map shows how the backbone circuit has changed out from 1998 gradually.

For a time, it was the case that circuits and router interfaces with quadruple the bandwidth were released roughly every two years. Bandwidth continued to expand to 600Mbps (STM-4) in 2000, 2.4Gbps (STM-16) in 2002, and 9.6Gbps (STM-64) in 2004. It was a happy era in which circuit capacity expanded as traffic grew. While the circumstances may have been different for larger ISPs, IIJ was mostly able to keep up with the growing traffic by upgrading circuits without significantly changing its network topology.

The following graph, based on the backbone map, plots total backbone bandwidth between Japan and the US along with the highest-capacity circuits that were available in the backbone. The circuit bandwidth steadily increased from 1998 up until when STM-64 rolled out in 2004. But after 2005 through 2014, over 9 years, circuit bandwidth above

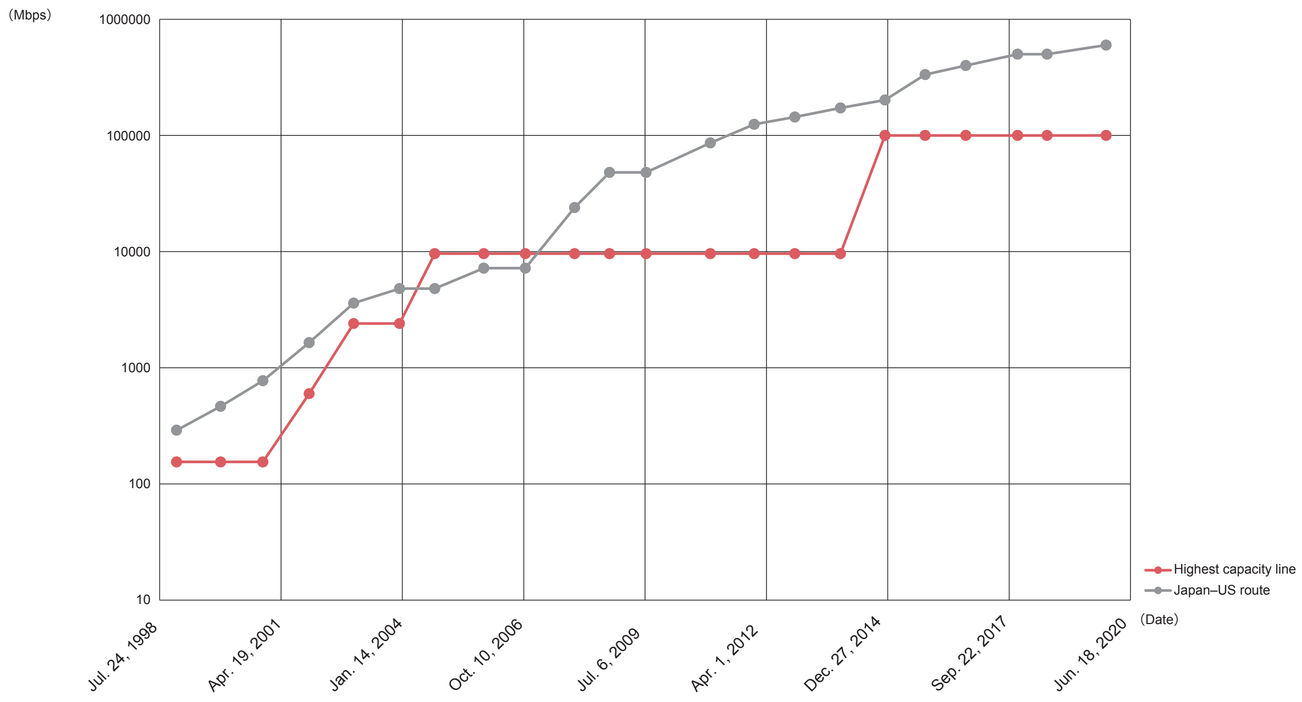

Figure 25: Line Capacity and Japan–US Bandwidth

STM-64 did not become standard. Various factors were behind this. STM-256 was not all that widely adopted, and technical difficulties associated with 100G, the next step up, caused development and standardization delays. Yet traffic continued to grow over this period, eventually reaching the point that we needed bandwidth of 200Gbps or 20 STM-64 circuits between Japan and the US. (Note that the graph's y-axis is logarithmic!)

Hence, the design of the IIJ backbone changed considerably over this period. To change our backbone topology freely, we adopted the MPLS abstract layer, and to boost circuit utilization efficiency, we altered the network topology to make it more amenable to ECMP (Equal Cost Multi-Path).

Things settled somewhat once 100G ethernet became standard in 2014. IIJ also finished converting core backbone circuits to 100G in 2018 and is in the process of upgrading the other network to 100G as traffic grows.

Meanwhile, a sixth 100G circuit between Japan and the US is in the schedule to go up in December 2019. With five years having passed since 100G went mainstream, we're starting to search for the next step up in capacity. The footsteps of 400G are drawing gradually closer.

IIJ will continue to follow the latest technological developments as it expands our backbone to ensure it remains fast and efficient with high availability.

1.BGP / Number of Route
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 FLET'S and Natural Disasters
Takafusa Hori
Manager, Network Technology Section, Network Technology Department, Infrastructure Engineering Division, IIJ

5.The Transition of the Bandwidth of IIJ Backbone Circuits
Takanori Sasai
Manager, Backbone Technology Section, Network Technology Department, Infrastructure Engineering Division, IIJ