Internet Trends as Seen from IIJ Infrastructure

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 and DNS query information, as well as IPv6 usage and mobile connectivity 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

Following on from last year's IIR Vol. 37 (https://www.iij.ad.jp/en/dev/iir/037.html), we start by looking at IPv4 full-route information advertised by our network to other organizations (Table 1, Figure 1). During the past year, RIPE NCC finished allocating/assigning its last /8 block (making it the second organization to do so, after ARIN). The size of allocations from the IANA Recovered IPv4 Pool to RIRs has also fallen to /22 (1,024 addresses). Increasingly, therefore, the acquisition of IPv4 addresses is becoming reliant on address transfers.

The last eight years has seen the largest increase in the total number of routes, which now exceeds 700,000. The growth rates for the /22 and /23 prefixes are above 10%, and taken together, the /22, /23, and /24 prefixes combined saw 89% growth in the number of routes to now account for 79% of all routes. As address blocks are increasingly split up for the purpose of transfer, it will be worth keeping an eye on the extent to which the proportion of routes accounted for by these prefixes grows.

Next we take a look at IPv6 full-route data (Table 2). The biggest increase here has also come in the past eight years.

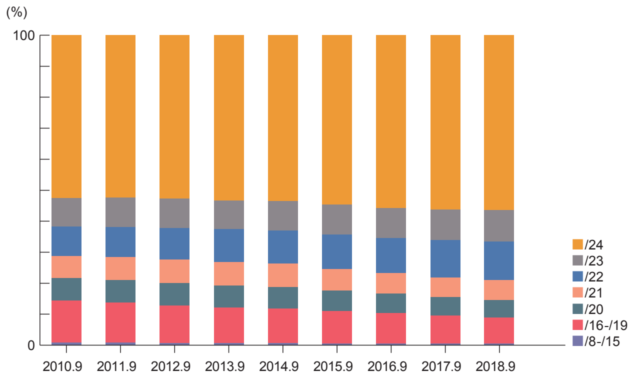

Figure 1: Percentage Breakdown of Number of Routes by Prefix Length for Full IPv4 Routes

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

	/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
Growth rate*	-1	-2	0	-5	8	15	47	30	-66	287	386	635	704	3948	9697	7481	33014	56178

*Since September 2017

That said, prefixes /33 through /48 account for 74% of all routes, and of those routes, we calculate that over 76% correspond to route advertisements for blocks that have been split into smaller fragments since being allocated (assigned). The number of routes is one gauge of the spread of IPv6, so an increase in this metric is desirable, but with fragments accounting for over half of the total, we seem to be somewhat removed from the original IPv6 ideal (?) of being able to limit growth of the routing table through consolidated route advertising, and this is a little disappointing.

Lastly, let's also take a look at IPv4/IPv6 full-route Origin AS figures (Table 3). With IANA's 16-bit Autonomous System Number (ASN) Pool having been exhausted in July 2016, the number of 16-bit Origin ASNs turned downward from 2016. The number of 32-bit only ASNs (allocation started in January 2007), meanwhile, has continued to rise steadily, but the vast majority appear to be operated only in the IPv4 space. This seems to indicate that even new organizations have not given much thought to using IPv6 despite probably having acquired ASNs and started running BGP during a time when the stockpile of IPv4 addresses had been exhausted, and this tells us that the IPv6 rollout still has a long way to go. That said, IPv4 addresses will certainly be harder and harder to obtain ahead, so whether this trends persists or not is also something we will be keeping close tabs on.

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

	/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
Growth rate*	10	522	72	1808	1240	823	326	2016	287	6269	13373

*Since September 2017

Table 3: 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%)

Topic 2

DNS

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 May 7, 2018.

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 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 4:00 a.m., and a peak of about 0.22 queries/sec per source IP address at around 1:00 p.m. Broken down by protocol (IPv4 and IPv6), the trends in query volume are virtually the same (no major differences) during daytime hours, whereas IPv6 queries per IP address show a tendency to rise after 8:00 p.m. This suggests that the computing environment needed to allow use of IPv6 in the home is coming into place.

Recent years have seen a tendency for queries to rise briefly at certain round-number times, such as hour marks. The number of query sources also increases, which tells us that the increase is possibly due to tasks being scheduled on user devices and automated network access that occurs when a device is activated by, for example, an alarm clock function. Diving a little deeper, we note an increase in queries 14 seconds before every hour mark. The increase in queries that occurs on the hour tapers of gradually, but with the spike that occurs 14 seconds before the hour, query volume immediately returns to about where it was. Hence, because a large number of devices are sending queries in almost perfect sync, we can infer that some sort of lightweight, quickly completed tasks are being executed. Some implementations, for example, may have a mechanism for completing basic tasks, such as connectivity tests or time synchronization, before bringing the device fully out of sleep mode, and the queries used for these tasks could be behind the spike.

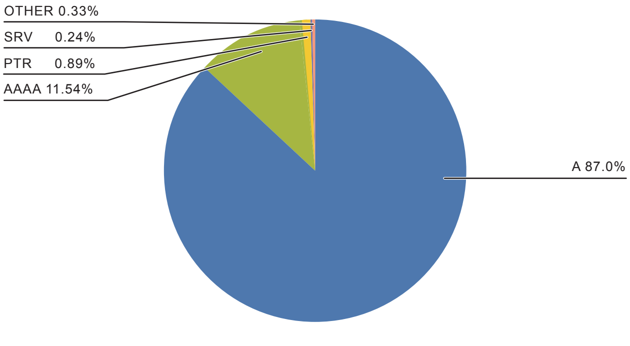

Figure 2: IPv4-based Queries from Clients

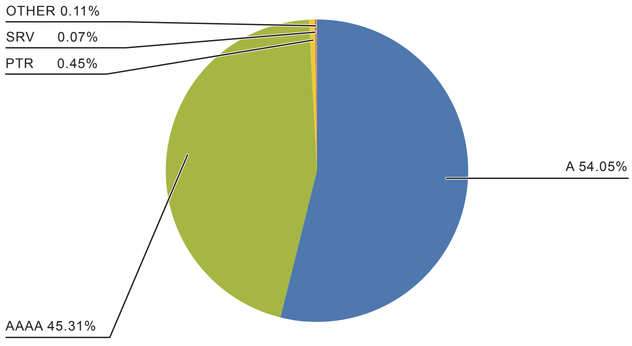

Figure 3: IPv6-based Queries from Clients

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. ANY queries have decreased relative to last year.

With ANY queries being widely abused for reflection attacks and the IETF continuing to discuss relevant countermeasures, this type of query is gradually falling out of use, which would explain the decline this year. Turning to trends broken down by query IP protocol, the number of query source IPs and the number of actual queries are both higher for IPv6-based queries than for IPv4 queries. 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 87% are A record queries and 11% AAAA record queries (Figure 2). With IPv6-based queries, meanwhile, AAAA record queries account for a higher share of the total, with around 54% being A record and 45% being AAAA record queries (Figure 3).

Topic 3

IPv6

Around a year has passed since we last looked at the state of IPv6 in Internet Infrastructure Review Vol. 37. In this issue, we look at what volume of overall traffic on the IIJ backbone is IPv6 and what protocols are mainly being used. We also look specifically at the state of and factors behind mobile services traffic, an area where IPv6 traffic is on the rise.

■ 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 4. The data span the year from October 1, 2017 to September 30, 2018. Over the year, IPv4 traffic increased by around 20% while IPv6 traffic rose by around 80%. IPv6 accounts for around 6% of overall traffic, an increase from around 4% last year. Figure 5 plots the data for the same period on a

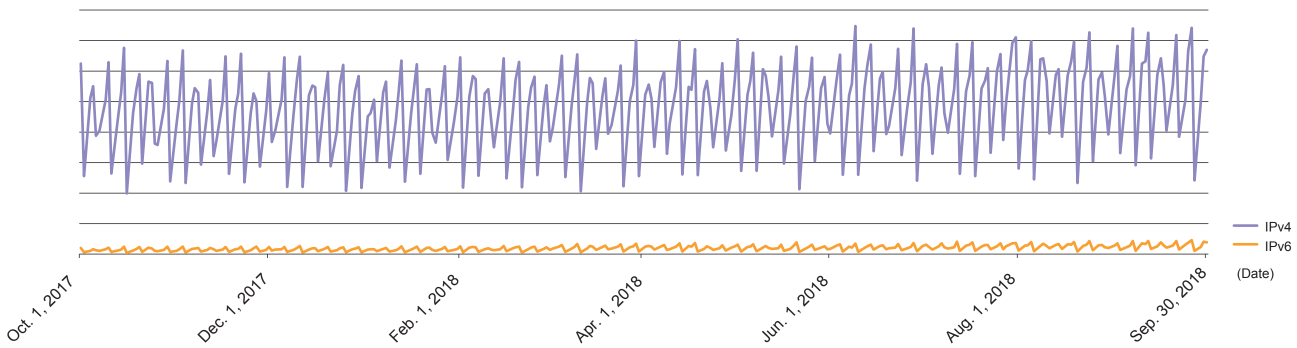

Figure 4: IPv4/IPv6 Traffic Measured via IIJ Backbone Routers at Core Points of Presence (Tokyo, Osaka, and Nagoya)

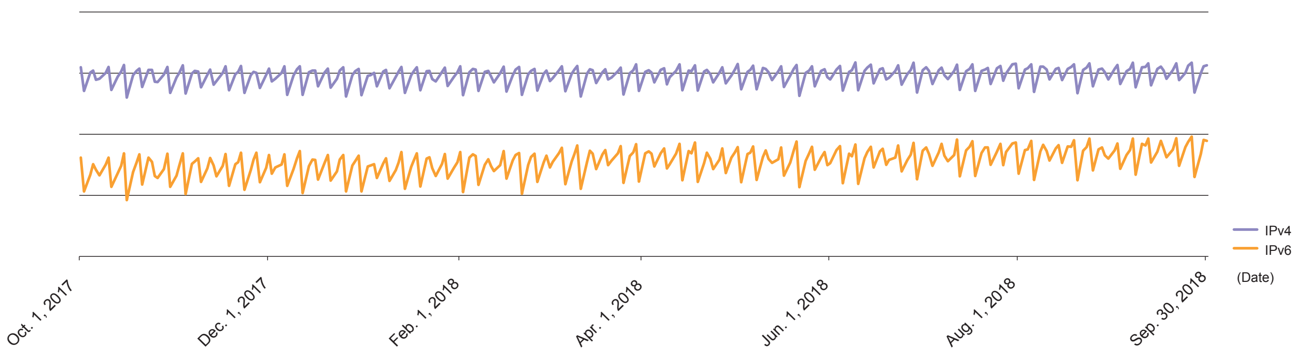

Figure 5: IPv4/IPv6 Traffic Measured via IIJ Backbone Routers at Core Points of Presence (Tokyo, Osaka, and Nagoya)—Log Scale

log scale. Although the absolute volume of traffic accounted for by IPv6 is less than a tenth that for IPv4, the rate of growth for IPv6 is clearly higher than for IPv4.

Next, Figures 6 and 7 show the top annual average IPv6 and IPv4 traffic source organizations (BGP AS Number) for the year from October 2017 through September 2018.

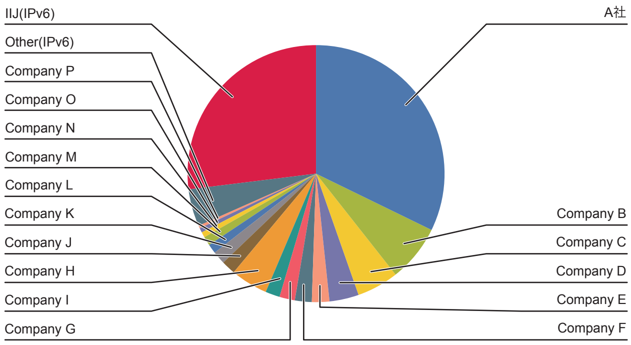

Figure 6: Top Annual Average IPv6 Traffic Source Organizations (BGP AS Number) from October 2017 to September 2018

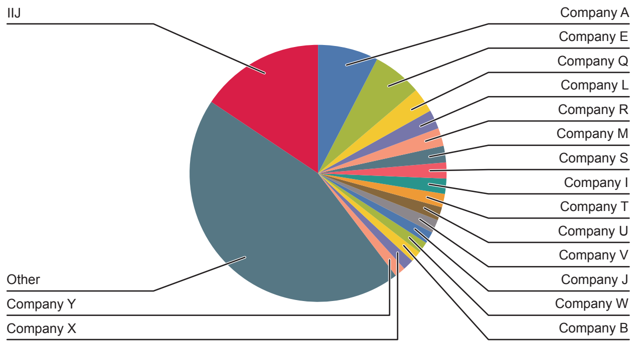

Figure 7: Top Annual Average IPv4 Traffic Source Organizations (BGP AS Number) from October 2017 to September 2018

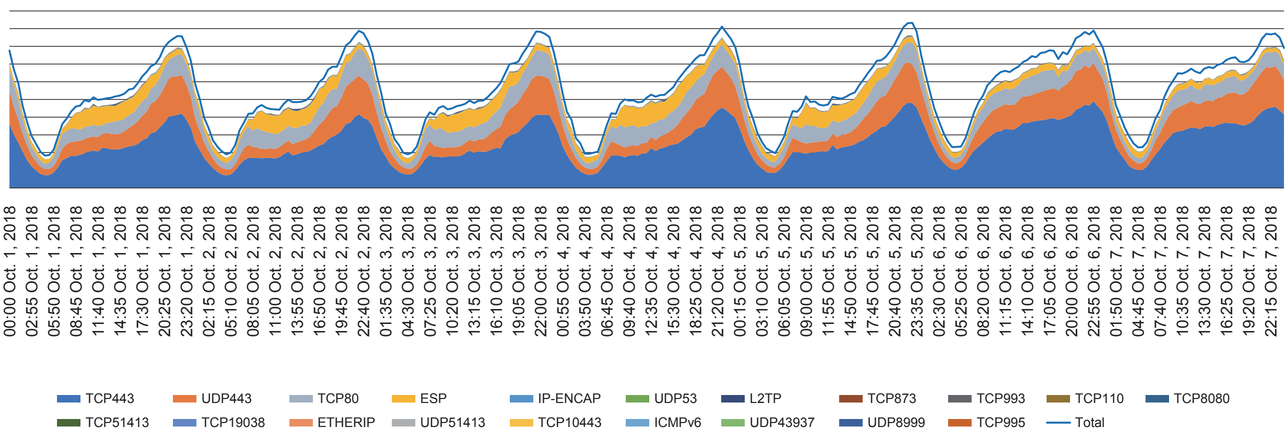

Figure 8: Breakdown of IPv6 Traffic by Protocol Number (Next Header) and Source Port Number

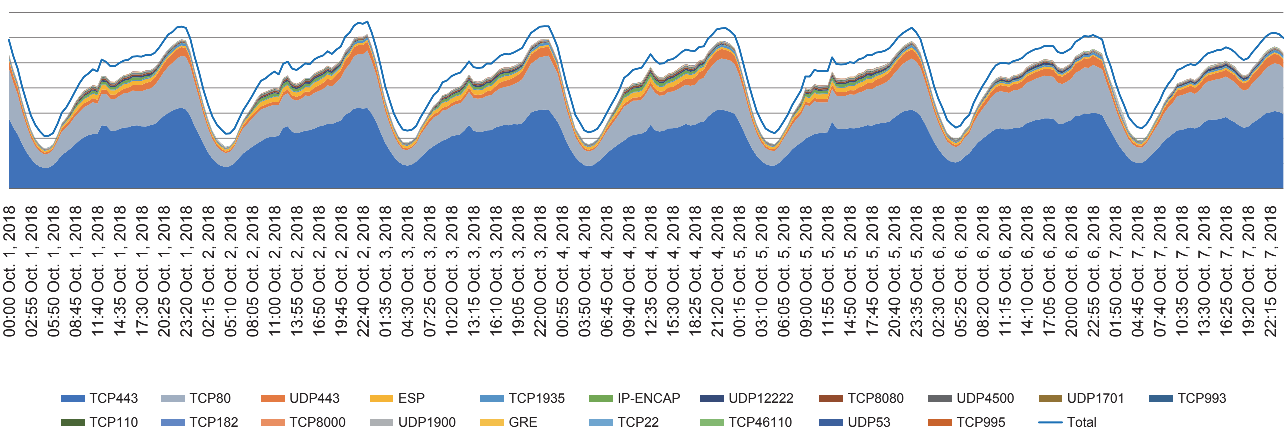

Figure 9: Breakdown of IPv4 Traffic by Protocol Number and Source Port Number

Last year's top ranking organization is still in the No. 1 spot but with its slice of the pie now only about half what it was due to a narrowing of the traffic volume gap versus No. 2 down. The data indicate that use of IPv6 is progressing at the No. 2 level on downwards as well. Providers of IPv6 IPoE access services via FLET'S Hikari Next come in at No. 4 and No. 6, which may tell us that the spread of IPv6 IPoE is leading to an increase in the use of IPv6.

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 October 1, 2018).

In a trend similar to last year, TCP/UDP 443 and TCP 80 now account for an even greater share of the total, with Web-based applications accounting for most of the traffic. This goes for not only IPv6 but IPv4 as well.

In addition, IP-ENCAP (Protocol Number 4) rose from No. 6 last year to No. 5 in the rankings this year. Although the series is too thin to be visible in the plots, the traffic numbers have more than doubled since last year, which we surmise indicates an increase in traffic using IPv4-over-IPv6 technologies such as DS-Lite (RFC6333).

Mobile Services IPv6 Traffic

In a new addition to this periodic report, we now look at IPv6 traffic on mobile services.

Figure 10 plots traffic on IIJ mobile services over a two-year period from October 1, 2016. IPv6 traffic starts to surge right around the middle of the plot, corresponding to late September 2017.

This coincides with the release of iOS version 11, the operating system on US-company Apple's iPhones and iPads. In iOS 11, the MNVO APN profile (config file for mobile

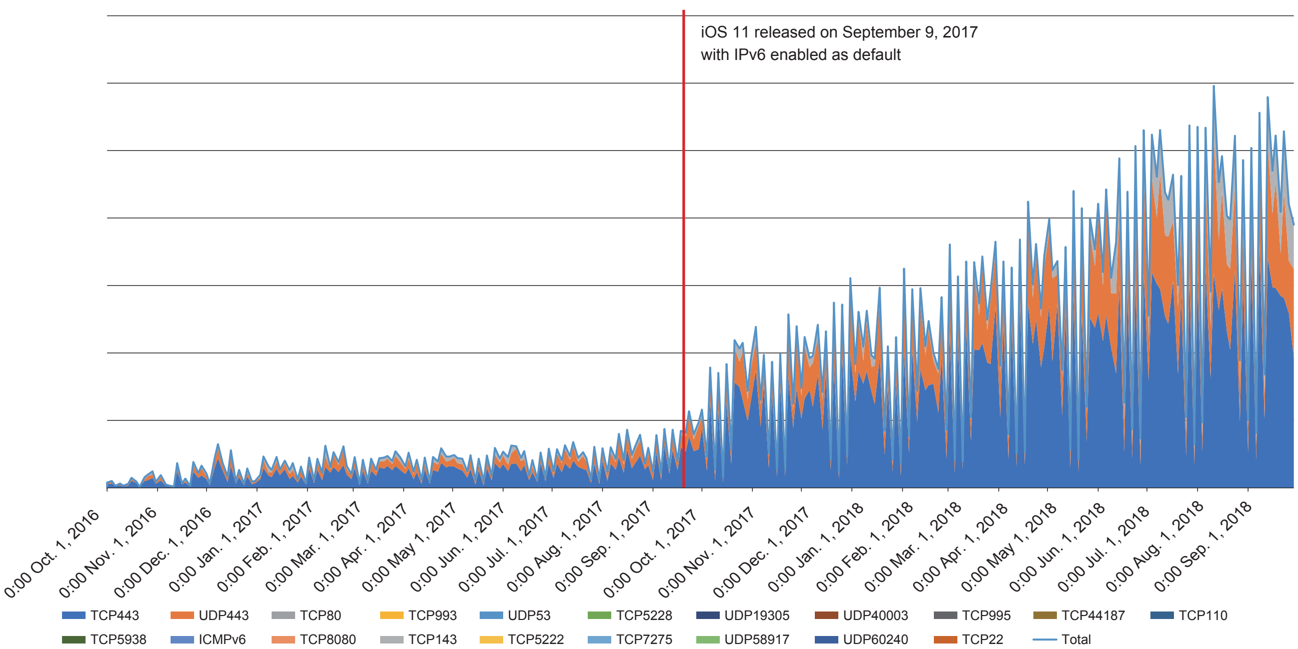

Figure 10: Mobile Services Traffic over Two Years from October 1, 2016

network connections) defaults to IPv6, and this may be why a lot of user devices started using IPv6.

Web-based application traffic accounts for almost all (over 98%) of mobile IPv6 traffic.

Summary

In this issue, we examined IPv6 traffic volume, protocols used, and IPv6 traffic for mobile services as a separate category. Overall, IPv6 traffic is on the rise, with the rate of growth outpacing IPv4. One reason for this seems to be the advance of IPv6 on everyday user devices due to factors such as the spread of IPv6 IPoE services on FLET'S Hikari Next and the release of Apple iOS 11, as well as the range of IPv6 -based service providers growing increasingly diverse. The last regional registry with IPv4 addresses remaining, AFRINIC (RIR for the African region), is expected to exhaust its IPv4 pool around the middle of 2019, so the use of IPv6 looks set to rise further and further.

Topic 4

Mobile and Broadband

We now analyze mobile and broadband traffic. Note that broadband in this section does not include FLET'S IPoE.

Figure 11 plots download and upload (from the user's perspective) traffic volume (bps) for both mobile and broadband normalized based on the peak values for each series. The plot shows that the peak in mobile traffic comes around noon and in broadband around 10:00 p.m. Mobile connections are often used when people are out of the house, so traffic is higher during the day. Increases can also be observed around times when people are commuting to and from work or school, so the data correlate strongly with people's daily movements. Broadband, meanwhile, is generally used at home once people return for the day, so traffic is heavier at night.

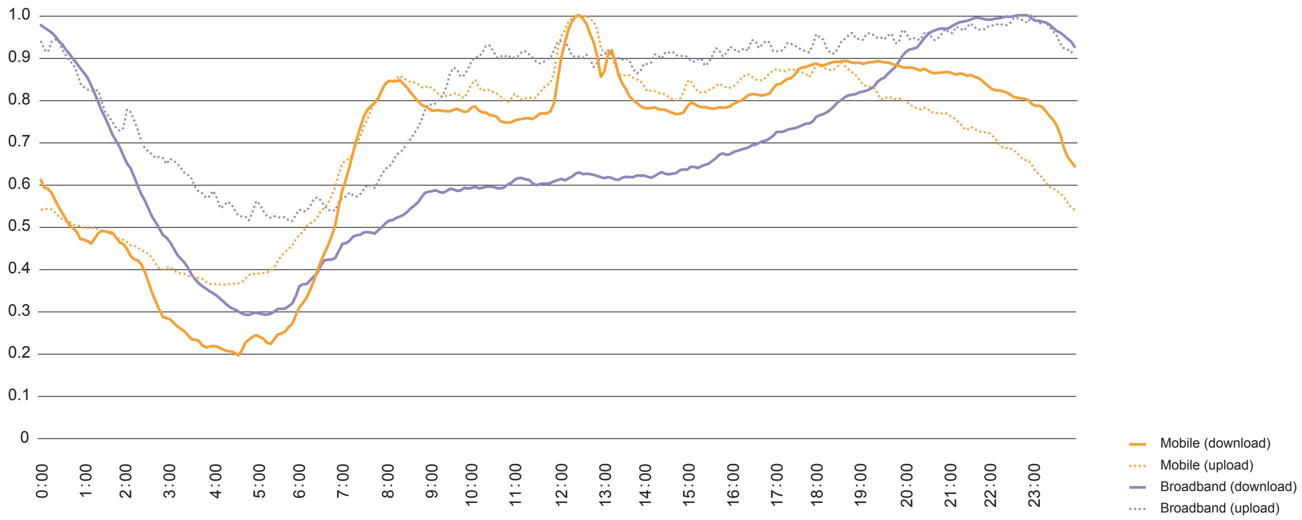

Figure 11: Traffic Indexed to Peak Levels

The fluctuations over the course of a day are larger for both mobile and broadband. Broadband traffic tends to rise gradually over the course of the day, from morning through to the nighttime peak. Mobile traffic, on the other hand, rises sharply in the morning and hovers at high levels up until right before midnight.

For reference, information on the traffic of Japan's five mobile carriers published in a Ministry of Internal Affairs and Communications document titled "Current State of Mobile Communications Traffic in Japan (June 2018)" [in Japanese] also indicates that traffic peaks at nighttime, as it does for IIJ broadband. Compared with MNO services, IIJ's mobile services (MVNO) currently attract more customers in the early adopter segment. We surmise that these customers also have broadband connections in the home and that they offload a high proportion of their nighttime traffic to broadband. As MVNO services spread further and attract more customers from the majority customer segments, the traffic peak is likely to shift into the nighttime, like that for MNO services.

Next, we compare download-to-upload ratios. Figure 12 shows this ratio, found by dividing download traffic (bps) by upload traffic, for both mobile and broadband.

The plot shows that the download ratio is higher for broadband than it is for mobile. Technological advances continue to push mobile communication speeds higher, but broadband still generally offers a more stable, high-speed connection. Plus, while broadband data is basically unlimited, mobile services are often subject to various forms of data transfer caps. This is probably why broadband sees higher capacity downloads.

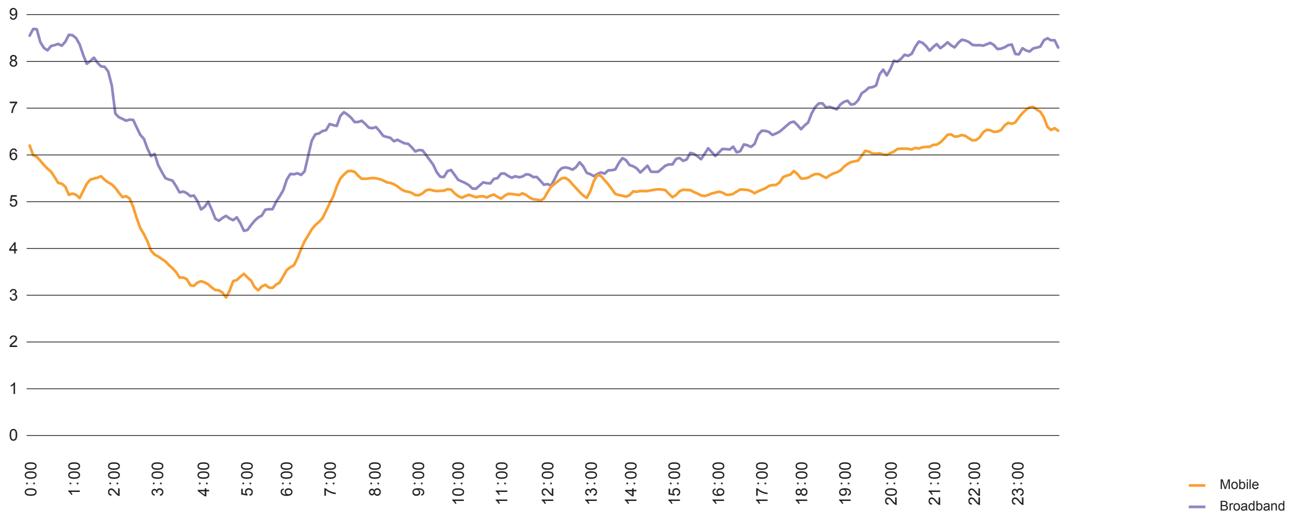

Figure 12: Download Relative to Upload Traffic

Next, we compare protocols. Figures 13 and 14 show a percentage breakdown (protocol, source port) of download traffic volume (bps) for mobile and broadband, respectively.

In both mobile and broadband, HTTP protocols (443/tcp and 80/tcp) account for about three quarters of the total. Also prominent is another relatively new protocol called QUIC, which uses 443/udp. The QUIC destinations are noticeably biased toward specific Internet service providers. Interestingly, 443/tcp (i.e., HTTPS) accounts for a greater share in mobile than in broadband. Most mobile users are quite possibly using smartphones, but rather than simply surfing Web pages via a stock-standard browser, it is possible that they frequently have occasion to use various applications designed for specific purposes and that the HTTPS protocol is commonly used by such applications.

Next, we look at IPv6 usage rates. Table 4 gives a percentage breakdown by connection type for mobile and broadband. The table shows that IPv6 usage is higher for mobile, albeit only slightly. With most new smartphones sold today being IPv6 capable, the groundwork for using IPv6 is naturally falling into place, even if users are unaware of it. With NTT's FLET'S service, users who have a compatible home gateway can use IPv6 without consciously having to do anything, but those users who provide their own broadband router need to configure it to use IPv6 themselves. FLET'S offers two connection types, PPPoE and IPoE, with the number of users connecting via IPoE, which faces fewer speed bottlenecks, rising of late. IPv6 is standard for IPoE connections, and to use IPv4 on IPoE connections, users need an environment that provides protocols to enable IPv4 over IPv6, such as DS-Lite.

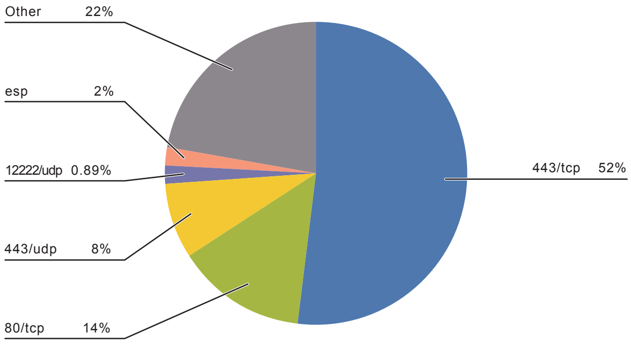

Figure 13: Percentage Breakdown of Download Traffic Volume (bps) by Protocol (Mobile)

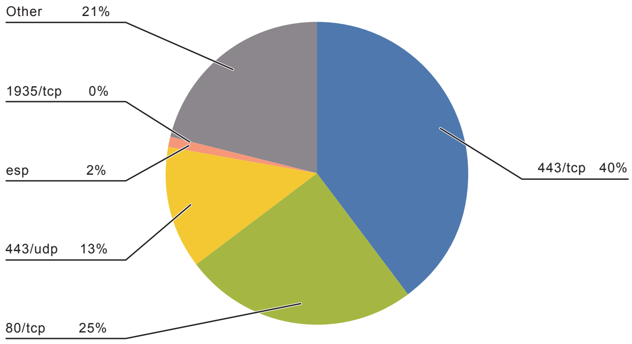

Figure 14: Percentage Breakdown of Download Traffic Volume (bps) by Protocol (Broadband)

A certain number of mobile users are exclusively IPv6. The number of services offering equivalent content via both IPv6 and IPv4 is rising, so it may be possible to meet demand in some cases with IPv6 alone, but it is surprising to see that there are so many of these users (when configured to use IPv4/IPv6 simultaneously, some devices make separate IPv4 and IPv6 connections for some reason, and this could be behind what we are observing).

Finally, let's break down mobile use by device. Mobile user communications use a protocol called GTP, and as part of the GTP connection process, devices report their IMEI (International Mobile Equipment Identity) to the network. IMEIs are composed of information that includes the device manufacturer and product name, so going through this data can give us a decent idea of what devices users are using.

Figure 15 shows a percentage breakdown by manufacturer. Even in a global context, Apple devices are noted as being highly prevalent in Japan, and they account for almost 40% of devices on IIJ's personal mobile services. This figure is astonishing given that, although IIJ does sell smartphones, it does not carry Apple products. And this shows just how appealing users find Apple devices to be.

The mobile user experience is very heavily influenced by the device (i.e., smartphone) used, and the mobile carriers put a lot of effort into smartphone sales for this reason. MVNO users, in some cases, continue to use the device they had when signing up for an MNO service, and so understanding what sort of smartphones are being used, including what devices are being brought over from MNO services, is crucial to the services strategy of MVNOs like IIJ.

Table 4: Breakdown of Connection Types for Mobile and Broadband

	IPv4	IPv6	IPv4v6
Mobile	70.43%	0.02%	29.55%
Broadband	75.48%	24.52%	NA

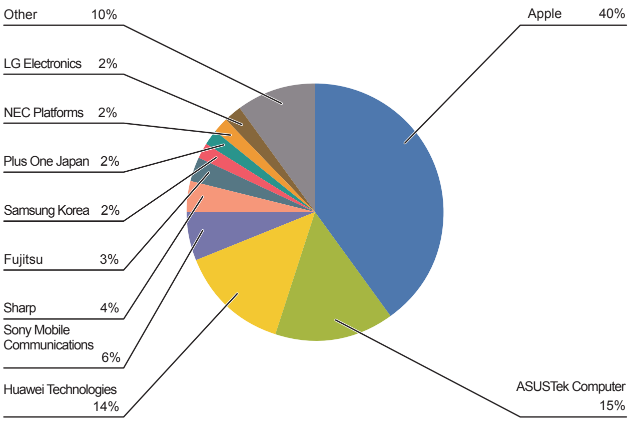

Figure 15: Device Breakdown

Topic 5

IIJ Infrastructure (Backbone)

IIJ monitors network status from a variety of angles to enable it to operate the IIJ network properly. In this issue, we take a look at one of the key metrics used: total traffic.

IIJ is an ISP, and one metric used to gauge an ISP's size is total traffic. Yet we cannot find many instances of what total traffic actually indicates being clearly explained. Here, we define total traffic to be the total bandwidth of transmissions into and out of IIJ's backbone. Backbone here means our routers collectively. It does not include transmission source/destination hosts. Although some transmissions are directed at the routers themselves, the volume is negligible.

With these definitions in place, let's look at IIJ backbone inflows/outflows. We can break these into three broad categories.

1. IIJ connection services customers
   - Connection services including Internet access services and datacenter access services
   - IIJ GIO (cloud service) Internet access service
   - Broadband (NTT East and West's FLET'S etc.) access services
   - Mobile connectivity services
2. IIJ service hosts
   - Email- and Web-related services
   - Content delivery
3. Interconnection business
   - Interconnections with other ISPs (peer)
   - Interconnections with cloud/content businesses

Based on this, Figure 16 shows plots of total traffic of the past 10 years. The data series are stacked. The outbound data are observations made at entry points, and the inbound data are observations made at exit points. Some traffic is eliminated within the backbone, such as that involved in

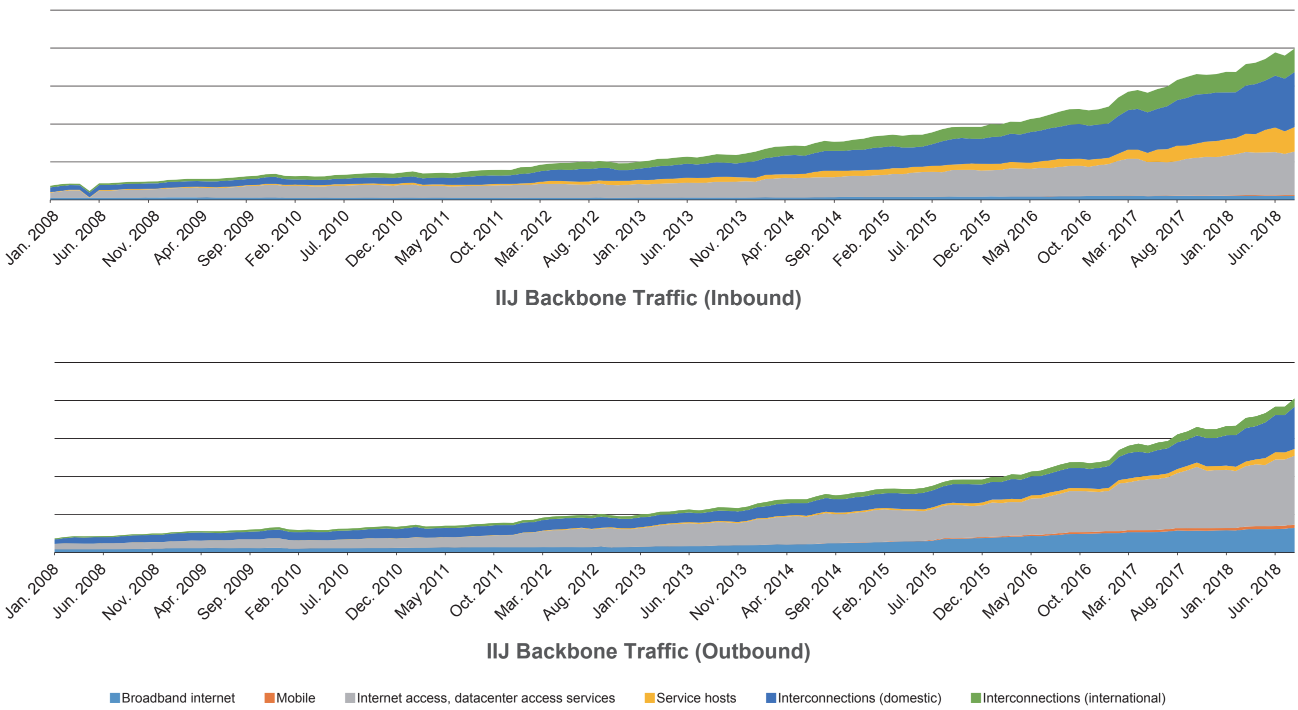

Figure 16: Total Traffic over 2008–2018

attacks, but generally all traffic that comes into the backbone also exits at some point, so the totals are almost the same.

The first thing you notice is that traffic has grown more than 10-fold over the past 10 years, with that growth accelerating. It shows no signs of easing. In the outbound plot, the three bottom series represent traffic that goes out to customers. Mobile traffic has been rising steadily since about three years ago but still accounts for only a small proportion of the total. Roughly speaking, the bulk of broadband and mobile traffic is accounted for by personal-use customers. Although this is rising steadily, the rate of growth in traffic to business customers (includes business customers to which IIJ provides personal services), shown in gray, has been higher over the past 10 years. Among IIJ customers, business traffic is growing more than personal-use traffic.

Let's look at inbound traffic. There tends to be less broadband and mobile traffic here, which means that not much information is sent out. Broadband traffic is growing, but the rate of growth is small, so it is accounting for a smaller and smaller proportion of the total. The share of traffic accounted for by service hosts is rising. This can be ascribed to the growth of content delivery and Web-based services.

Turning to a comparison of outbound and inbound traffic, we see that for broadband, although inbound traffic has only grown about twofold, outbound has grown close to ninefold. Here, too, it is evident that individual content offerings are becoming larger and larger. On the interconnection front, domestic inbound and outbound traffic look fairly well matched, whereas inbound is clearly greater for international interconnections. This seems to indicate that IIJ's content is not so much attractive when it comes to the international interconnection business.

This section has walked through a number of plots of IIJ's total traffic. Different insights emerge even from within the same traffic dataset depending on how you look at the data or what you focus on. Aside from traffic, we also record other completely different metrics such as latency or errors within the backbone. Going forward, we will continue to monitor IIJ's network and periodically report on changes in our observations.

1.BGP / Number of Route
Tomohiko Kurahashi
Infrastructure Planning Department, Service Infrastructure Division, IIJ.

2.DNS
Yoshinobu Matsuzaki
Infrastructure Planning Department, Service Infrastructure Division, IIJ

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

4.Mobile
Takafusa Hori
Manager, Network Technology Department, Service Infrastructure Division, IIJ

5.IIJ Infrastructure (Backbone)
Takanori Sasai
Manager, Backbone Technology Section, Network Technology Department, Service Infrastructure Division, IIJ