1. Periodic Observation Report

Internet Trends as Seen from IIJ Infrastructure

IIJ provides Internet services by operating some of the largest network and server infrastructure in Japan. Here, we examine and discuss current Internet trends based on information obtained by operating this infrastructure.

We cover the topics of network routing information and DNS query information, as well as the usage status of IPv6 and mobile access 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

Six and a half years have passed since February 3, 2011, when the pool of IPv4 addresses was exhausted at IANA, which oversee the allocation of global IP address resources. Currently, all five of the global RIRs (Regional Internet Registries) allocated IP addresses by IANA to allocate to individual countries have begun (or have already finished) the allocation or assignment of addresses from their last /8 block. Meanwhile, the number of IPv4 "full routes" observed on the Internet has been climbing steadily even after the exhaustion of IANA addresses, and is now close to double the amount as it was in 2011. In this section, we once again confirm trends in the number of routes based on IPv4 full-route equivalent information advertised from our network to other organizations (Table 1, Figure 1).

In general, the growth rate tends to be higher for longer prefix routes, which is what we expect. The notable increase in /22 routes is probably due to the size of IPv4 address space allocated/assigned by RIRs whose /8 block is almost depleted being limited to a maximum of /22 (1024 addresses).

Table 1: Trends in the Number of IPv4 Routes for Each Prefix Length Included in Full Routes

	/8	/9	/10	/11	/12	/13	/14	/15	/16	/17	/18	/19	/20	/21	/22	/23	/24	total
September 2010	20	10	25	67	198	409	718	1308	11225	5389	9225	18532	23267	23380	30451	29811	170701	324736
September 2011	19	12	27	81	233	457	794	1407	11909	5907	9885	19515	26476	26588	35515	34061	190276	363162
September 2012	19	14	29	84	236	471	838	1526	12334	6349	10710	20927	30049	31793	42007	39517	219343	416246
September 2013	16	11	30	93	250	480	903	1613	12748	6652	10971	22588	32202	34900	48915	42440	244822	459634
September 2014	16	12	30	90	261	500	983	1702	13009	7013	11659	24527	35175	37560	54065	47372	268660	502634
September 2015	18	13	36	96	261	500	999	1731	12863	7190	12317	25485	35904	38572	60900	52904	301381	551170
September 2016	16	13	36	101	267	515	1050	1767	13106	7782	12917	25229	38459	40066	67270	58965	335884	603443
September 2017	15	13	36	104	284	552	1047	1861	13391	7619	13385	24672	38704	41630	78779	64549	367474	654115
Growth Rate*	0.75	1.3	1.44	1.552	1.434	1.35	1.458	1.423	1.193	1.414	1.451	1.331	1.663	1.781	2.587	2.165	2.153	2.014

*September 2017 values with September 2010 value normalized to 1

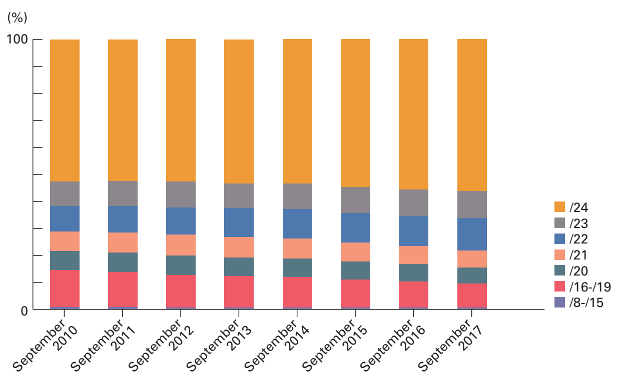

Figure 1: Trends in the Number of Routes for Each Prefix Length as a Proportion of IPv4 Full Routes

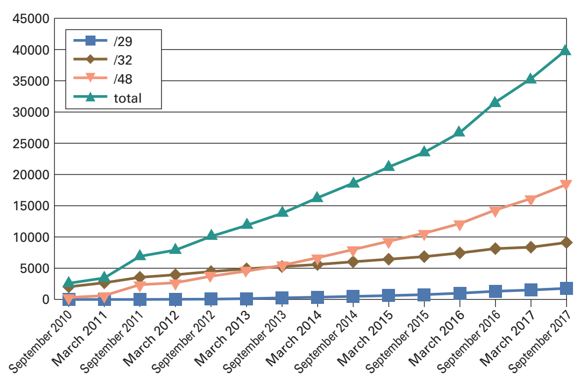

Figure 2: Trends in the Number of IPv6 Full Routes

You can also see that while the number of /8 routes is decreasing, all the others are increasing. We believe this is caused by the splitting of address blocks for the purpose of address transfers. Put simply, when an address block is split, the number of routes on the left side of Table 1 decreases, while the number of routes on the right side increases. It is likely that address transfers will continue to be used as a valuable means of obtaining IPv4 addresses, so we expect the distribution of the number of IPv4 routes to become even more biased toward the right side of Table 1 (with the longer prefixes).

In closing, we will touch upon the number of full routes for IPv6, the successor to IPv4 (Figure 2). Although the number is still insignificant compared to IPv4, it is increasing steadily, and from around 2016 this trend seems to have gained momentum. With the supply of IPv4 addresses from each RIR becoming further depleted, we speculate that support for IPv6 on a regional and organizational level will accelerate in the future. It will be interesting to see how this trend changes going forward.

Topic 2

DNS

IIJ provides a full resolver to enable users of our Internet access services to use DNS name resolution. In this section, we discuss the state of name resolution, and analyze and reflect upon data for servers mainly provided for broadband, based on a day's worth of full resolver observation data obtained by IIJ on May 17, 2017.

Full resolvers only know the IP address of the authoritative name server that provides top-level zone information, which is known as the root. Based on the information gained from there, they track down authoritative name servers likely to possess information for locating the required record. Because load and latency become an issue when making recursive queries each time using a full resolver, the records obtained are cached for a while, and retrieved from the cache when the same query is received again. Recently, functions related to DNS have also begun to be implemented in devices on the communication route, such as broadband routers or firewalls. These may be involved in relaying DNS queries or applying control policies.

On broadband and mobile connections, it is possible to use protocols such as PPP, DHCP, RA, and PCO to notify users of the IP address of the full resolver. ISPs use these functions to enable automatic configuration of full resolvers for the name resolution required by user communications. ISPs can inform users about multiple full resolvers, and users can specify and add full resolvers to use by changing settings themselves. When more than one full resolver is configured on a computer, the one used depends on the computer's implementation or application, so full resolvers are not aware of the total number of queries being sent by a user. This means that full resolvers must be operated with extra processing power in reserve, while keeping a close watch on query trends.

Looking at user trends in the observation data for the full resolver provided by IIJ, we observe a daily average of about 0.08 queries/sec per source IP address. This value fluctuates depending on the time of day, indicating trends in user activity, with about 0.04 queries/sec at around 4:00 a.m., and about 0.13 queries/sec during the peak at around 9:00 p.m. Both the IPv6 and IPv4 IP protocols are used for query communications, and these exhibit almost identical trends, with queries via IPv6 showing slightly greater variance due to the time of day. These values have not undergone any significant change over the past few years, staying within a variation range of about 0.06 points. Variable elements include the number of full resolvers that can be used by a client, the caching function in the user environment, and the behavior of computers and applications. This means it is hard to anticipate future trends, making continuous observation necessary.

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. There are differences in trends between the IP protocols used for query communications, with more AAAA record queries seen for IPv6 queries. For IPv4 queries, about 64% of the total are A record queries, while about 33% are AAAA record queries (Figure 3). Meanwhile, about 56% of total IPv6 queries are A record queries, and a higher ratio of about 43% are AAAA record queries (Figure 4). Also, examining trends for each query source IP address, about 96% are attempts to search for A records, regardless of whether the query was made via IPv4 or IPv6. As for AAAA records, about 57% of query sources in IPv4 and about 80% in IPv6 are searches. The ratio of records that account for IPv4 queries has not changed a lot in the past few years, so we surmise that new implementations in recent years may prioritize using IPv6 for queries.

Topic 3

IPv6

On February 3, 2011, the IPv4 address pool of the APNIC RIR that manages IP address resources in the Asia-Pacific region was exhausted, and the new allocation of regular IPv4 addresses (assignment to regional management organizations and ISPs) ended in Japan. In other words, the stock of IPv4 addresses was depleted. About six and a half years have passed since then, but it still cannot be said that there has been an explosive spread in the use of its successor, IPv6.

Here, we will analyze IPv6 user numbers, traffic and usage protocols at IIJ and explain their current state.

■ Number of Users

IIJ started offering IPv6 PPPoE connections to customers using the FLET'S HIKARI NEXT service of NTT East and NTT West in June 2011. From July 2011, we also began providing IPv6 IPoE connections (in collaboration with affiliate INTERNET MULTIFEED CO.). In July 2015, we launched support for IPv6 PPPoE automatic connection from home gateways rented out by NTT East/ West, so that customers can use IPv6 connection without special settings. On our mobile services, we also provided support for IPv6 connections since the launch of our 4G (LTE) plans in May 2012. This enables mobile IPv6 connections as long as the device supports IPv6.

Figure 5 shows trends in the number of IPv4 and IPv6 connections on FLET'S HIKARI NEXT from July 2015 to the end of September 2017. IPv4 saw a negligible decrease, while IPv6 increased slightly, and as of September 2017 IPv6 accounts for approximately 22.9% of total connections (PPPoE 22%, IPoE 0.9%).

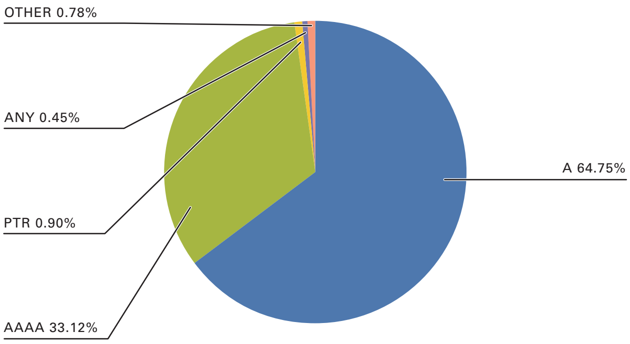

Figure 3: IPv4-based Queries From Clients

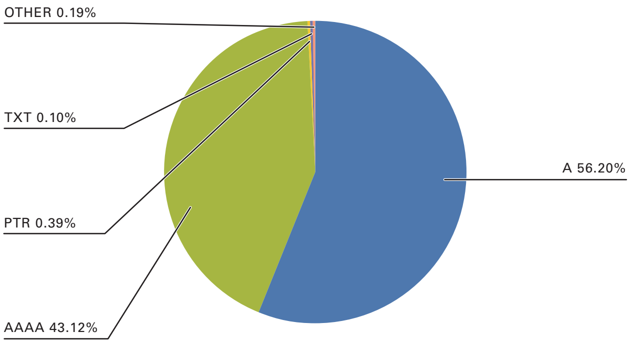

Figure 4: IPv6-based Queries From Clients

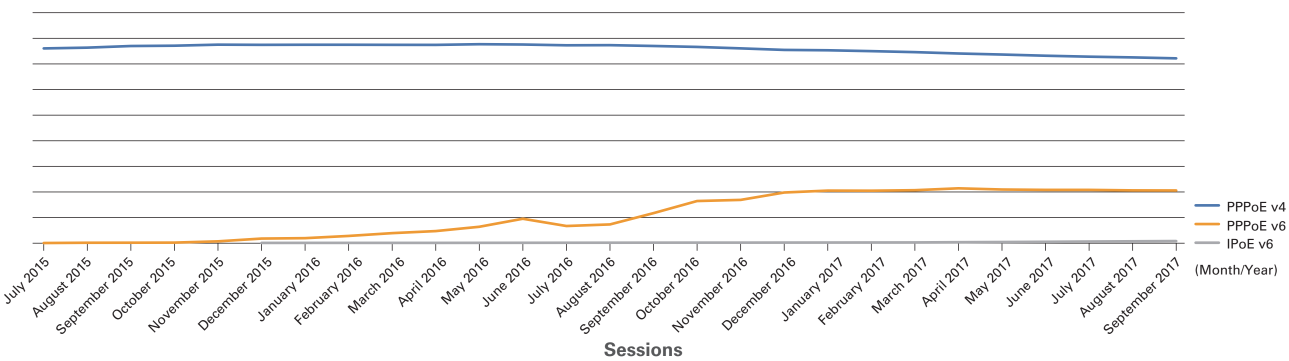

Figure 5: Trends in the Number of IPv4 and IPv6 Connections on FLET'S HIKARI NEXT from July 2015 to the End of September 2017

We believe that the number of IPv6 users has remained at around 22.9% because some users have devices that do not support IPv6 PPPoE automatic connections, and others have signed corporate contracts that do not provide IPv6, so we do not foresee there being any big increases in IPv6 PPPoE from here on. While IPv6 IPoE numbers are still low right now, we expect that migration to IPv6 IPoE using DS-Lite will gradually progress, and the gap will shrink.

■ Traffic

Figure 6 shows IPv4 and IPv6 traffic measured using IIJ backbone routers at core POPs (points of presence – in Tokyo, Osaka, and Nagoya). Both IPv4 and IPv6 are climbing, but IPv6 traffic only accounts for about 4% of the total. As a result, it is overshadowed by the IPv4 results when the two are put alongside, so at the moment it is difficult to say that it is gaining in popularity.

Next, Figure 7 shows the top annual average IPv6 traffic source organizations (BGP AS number) from October 2016 to September 2017.

At the top of the rankings is Company A, which is actively moving ahead with IPv6 support on its services, with the companies second and lower reaching just 1/16 of its numbers.

Company A also leads the IPv4 rankings (Figure 8), followed by major cloud vendor Company D in second, then major CDNs Company G and Company K. It is interesting to see the difference in types of organizations compared to IPv6. Another difference is that the number two company is at about 1/2 the level of Company A in the lead, unlike the differences for IPv6.

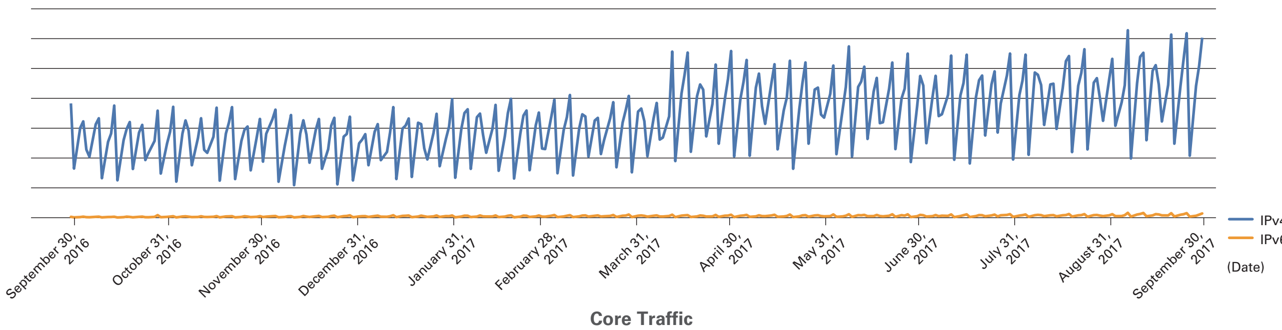

Figure 6: IPv4/IPv6 Traffic Measured Via IIJ Backbone Routers at Core Population Centers (Tokyo, Osaka, and Nagoya).

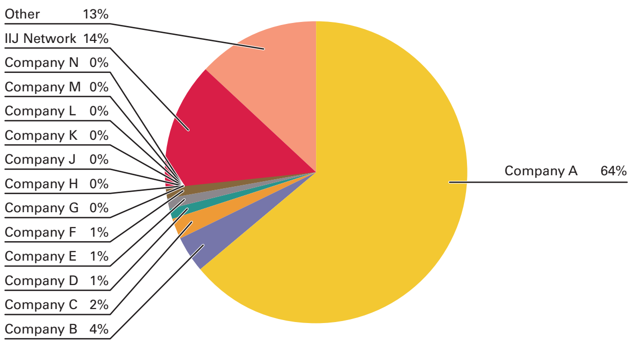

Figure 7: Top Annual Average IPv6 Traffic Source Organizations (BGP AS Number) from October 2016 to September 2017

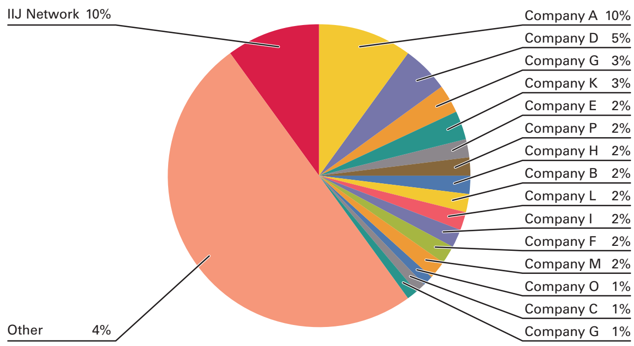

Figure 8: Top Annual Average IPv4 Traffic Source Organizations (BGP AS Number) from October 2016 to September 2017

Looking at the gap in traffic volume between Company A and those below, we would hazard a guess that although business operators other than Company A provide services with IPv6, its use may be limited.

■ Protocols Used

Figure 9 shows a graph analyzing via the protocol number (Next-Header) and source port number of IPv6 traffic (for a one-week period from October 1, 2017).

443/TCP (HTTPS) accounts for about 40% of the total, rising to over 50% when combined with 443/UDP (which we believe is QUIC) in second. Although 80/TCP (HTTP) is in third, this amounts to only about 1/6 of second and first combined, showing a difference that is quite remarkable compared with the IPv4 graph for the same period (Figure 10). We believe that IPv6 has a higher ratio of HTTPS/QUIC because Company A accounts for a large proportion of traffic. However, this may also be explained by new services from other companies that support IPv6 also being provided via HTTPS since its launch.

■ Summary

In this report we examined user numbers, traffic volumes, and usage protocols for the current state of IPv6 at IIJ. Although the environment for IPv6 connections has improved, the impression we get is that, aside from one company, support is only just beginning from the service provider side. With the three major Japanese carriers (NTT DOCOMO, KDDI, and SoftBank) having announced support for IPv6 one after another in 2017, we expect that support by service providers will gain momentum going forward. We will continue to analyze the situation from a variety of perspectives.

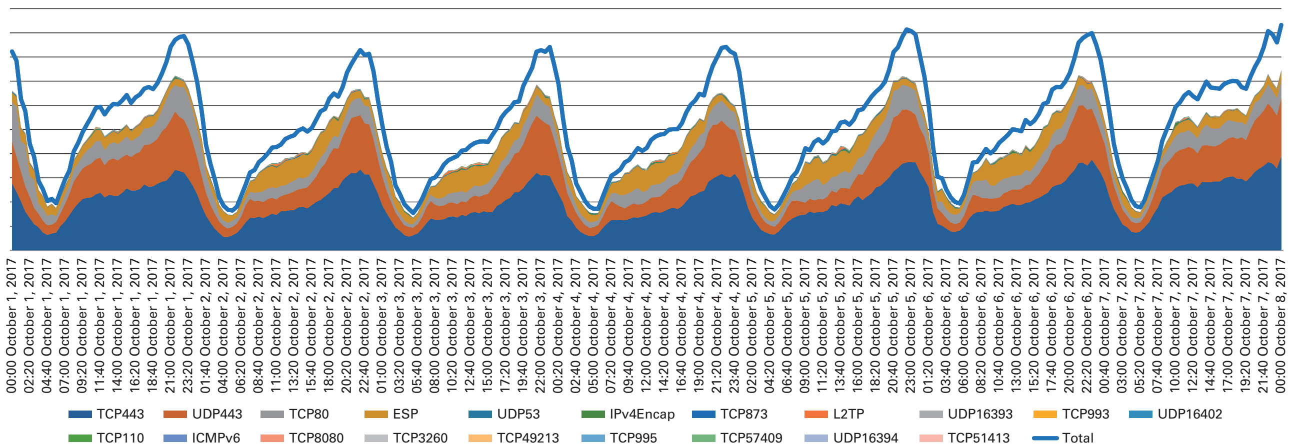

Figure 9: Graph Analyzing via Protocol Numbers (Next-Header) and Source Port Numbers of IPv6 Traffic

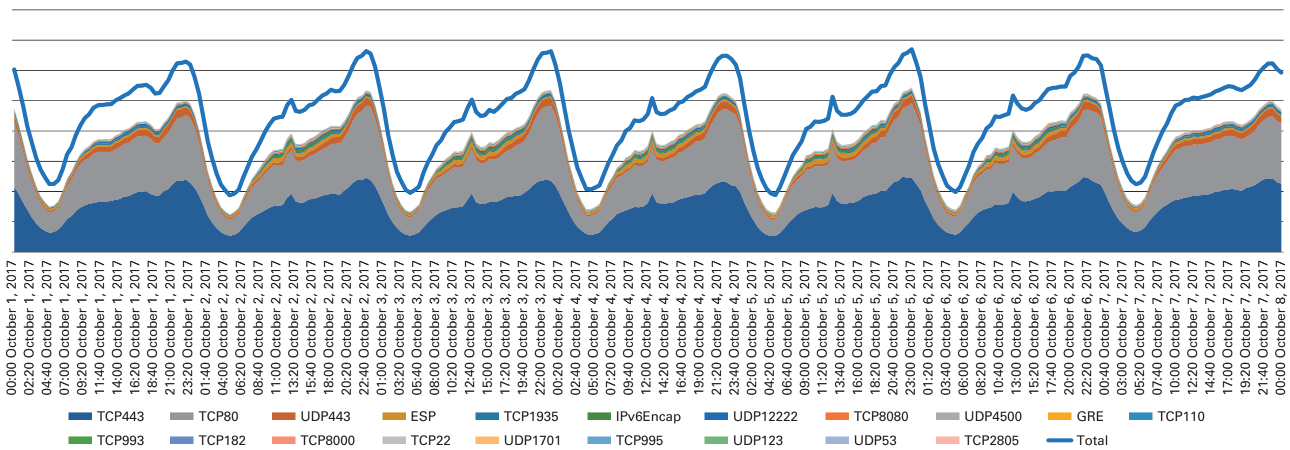

Figure 10: Graph Analyzing via Protocol Numbers and Source Port Numbers of IPv4 Traffic

Topic 4

Mobile

In this report, we analyze mobile traffic trends by focusing on time of day.

Figure 11 shows trends in traffic (bps) over the course of a single working weekday. The vertical axis represents relative traffic volume, while the graph shows the changing trends. These days, most mobile service users use smartphones. As you can imagine considering the situations where smartphones are used, the three large peaks in the graph correspond to the morning commute to work or school, the midday lunch break, and the commute back home from work or school in the evening. Traffic also drops sharply after 11:30 p.m.

We can see that usage is most concentrated at around 12:00 p.m. This is because, although the commute times to and from work and school in the morning and evening are spread out, lunch breaks are closely packed together around 12 o'clock. This causes congestion during this time of day. The TCP/IP mechanism controls traffic when congestion occurs, but despite this traffic levels remain high. To raise the utilization rate of facilities, it is important for ISPs to level out traffic variance at different times of the day by creating demand in areas other than smartphones, but this is not an easy task.

Figure 12 is a graph showing the traffic for a given week. You can see a similar pattern repeating from Monday to Friday. While the lunchtime peak at 12:00 p.m. is smaller on Saturday and Sunday, there is no decline in traffic during the day. The dip between Sunday night and Monday morning is also deeper. Although it is hard to see on this graph, the dip in traffic at night tends to become shallower over the weekend. This is an interesting trend that reflects our day-to-day activities.

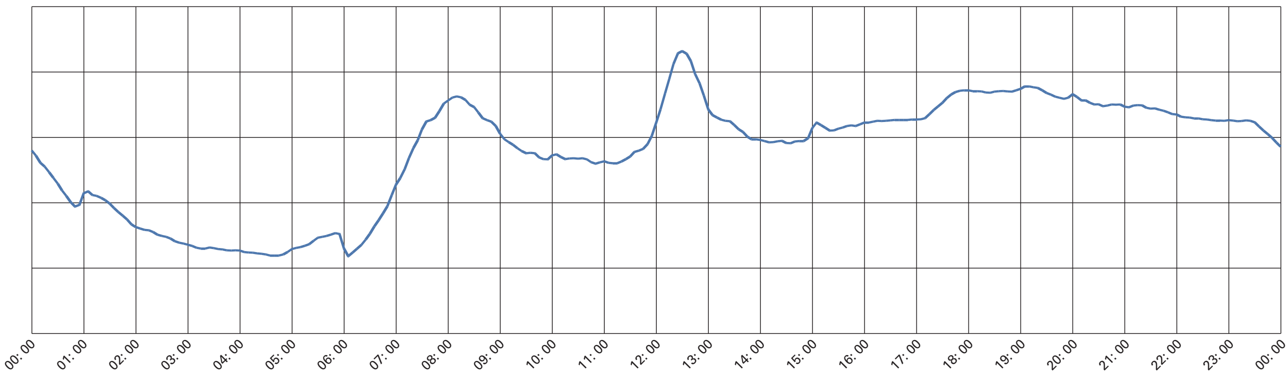

Figure 11: Download Traffic Trends for One Day

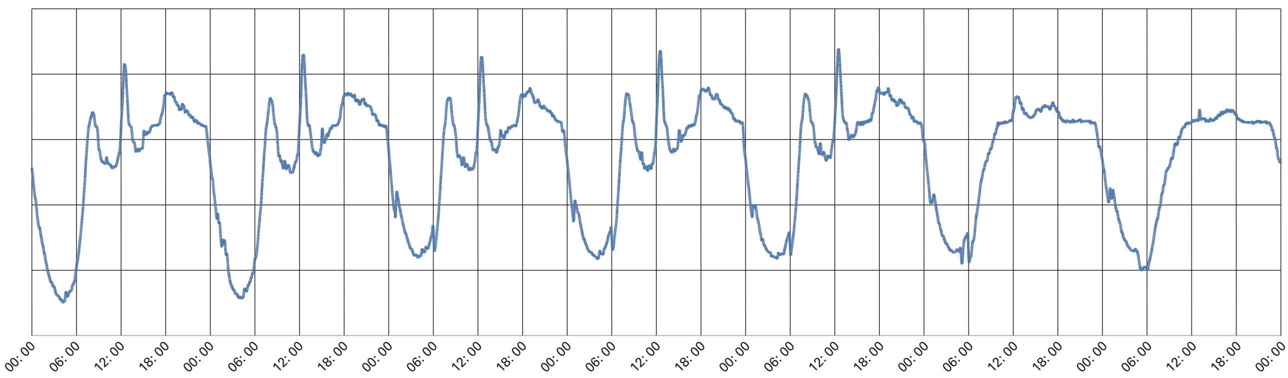

Figure 12: Download Traffic Trends for One Week

Figure 13 plots the daily data transfer volume calculated from the traffic graph for each day. This graph shows data for October, but almost identical trends can be seen in other months when there are no large consecutive holidays, such as the year-end and New Year holidays. Transfer volumes are small at the start of the week, but increase the closer you get to Friday, then decrease over the weekend. Transfer volumes may climb as the weekend approaches because more and more data is transferred at night. It is also thought that the transfer volumes go down on the weekend because data transfers are offloaded to broadband-connected devices in the home. Interestingly, transfer volumes also go down towards the end of the month. We think this is because communications are reduced for users who have used up their monthly data allowance, but we have not yet found evidence to support this. At the start of a new month, transfer volumes that had been low at the end of the previous month will recover to their original levels or higher. Looking at overall mobile traffic throughout the year, we can see it is increasing steadily.

For some people, smartphones are essential items closely related to our everyday lives. This can be seen clearly in mobile traffic trends.

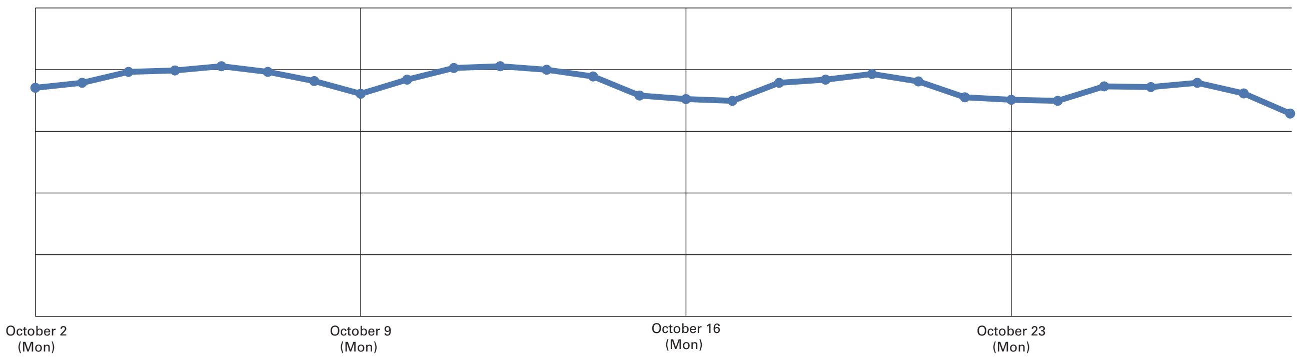

Figure 13: Data Transfer Volume by Date

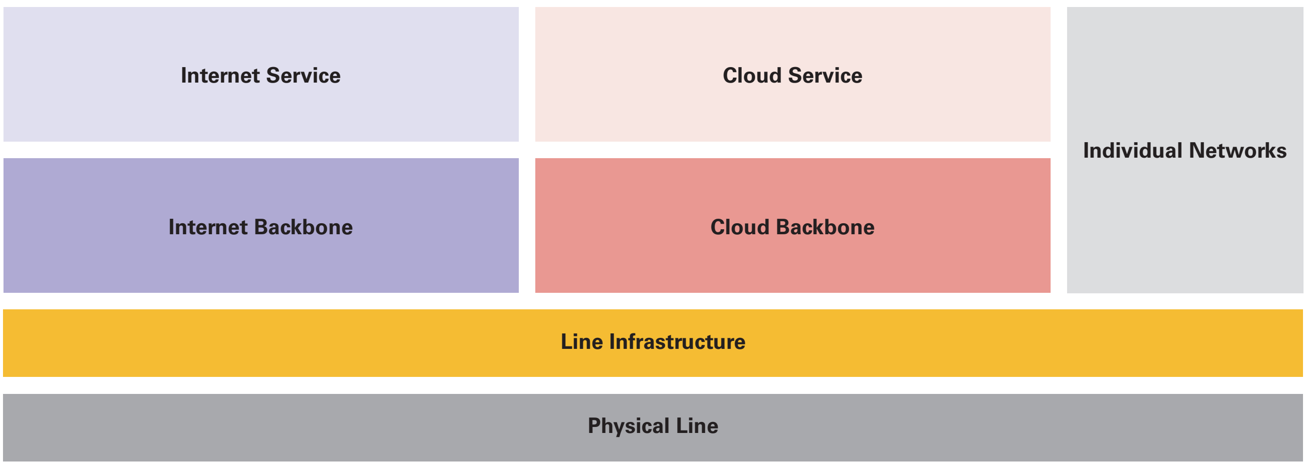

Figure 14: IIJ Backbone Structure

Topic 5

IIJ Infrastructure (Backbone)

Here we give an overview of IIJ backbone infrastructure.

Traffic is increasing at a steady rate. Overall Internet traffic has grown by a factor of 1.35 times per year on average, and the number of access lines in Japan and the U.S. has increased by 1.2 times on average annually (over the past four years in each case). Cloud traffic including services such as IIJGIO has also increased by 2.5 times over the past two and a half years.

Backbone infrastructure has evolved to support this increase in traffic. In terms of scale, we have extended the 100G line implemented between Tokyo, Nagoya, and Osaka three years ago to regional POPs, a connection between Japan and the U.S., and even to locations on the east coast of the U.S. Meanwhile, there have also been structural changes. Current backbone infrastructure establishes a layer 2 closed network (line infrastructure) for providing virtual lines over physical ones, and the backbone for Internet and cloud solutions is configured on the virtual lines provided by this infrastructure. Internet and cloud traffic are both provided over the same physical line, and traffic engineering implemented in the layer 2 closed network improves the efficiency of line utilization and increases the cost benefits. Another significant benefit is that this structural change has also made it possible to freely build networks without being bound by geographical restrictions. You could say the new IIJ DDoS Protection Service launched during the last fiscal year for handling large-scale attacks was made possible by this structural change. We will continue to evolve our backbone infrastructure to provide a variety of network services related to the Internet and cloud computing.

Authors:

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

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

3. IPv6
Taisuke Sasaki
Infrastructure Planning Department, Service Infrastructure Division, IIJ

4. Mobile
Takanori Sasai
Mobile Technology Section, Network Technology Department, Service Infrastructure Division, IIJ

5. IIJ Infrastructure (Backbone)
Daisuke Sugawara
Backbone Technology Section, Network Technology Department, Service Infrastructure Division, IIJ