IIR

Internet
Infrastructure
Review

Feb.2022
Vol.53

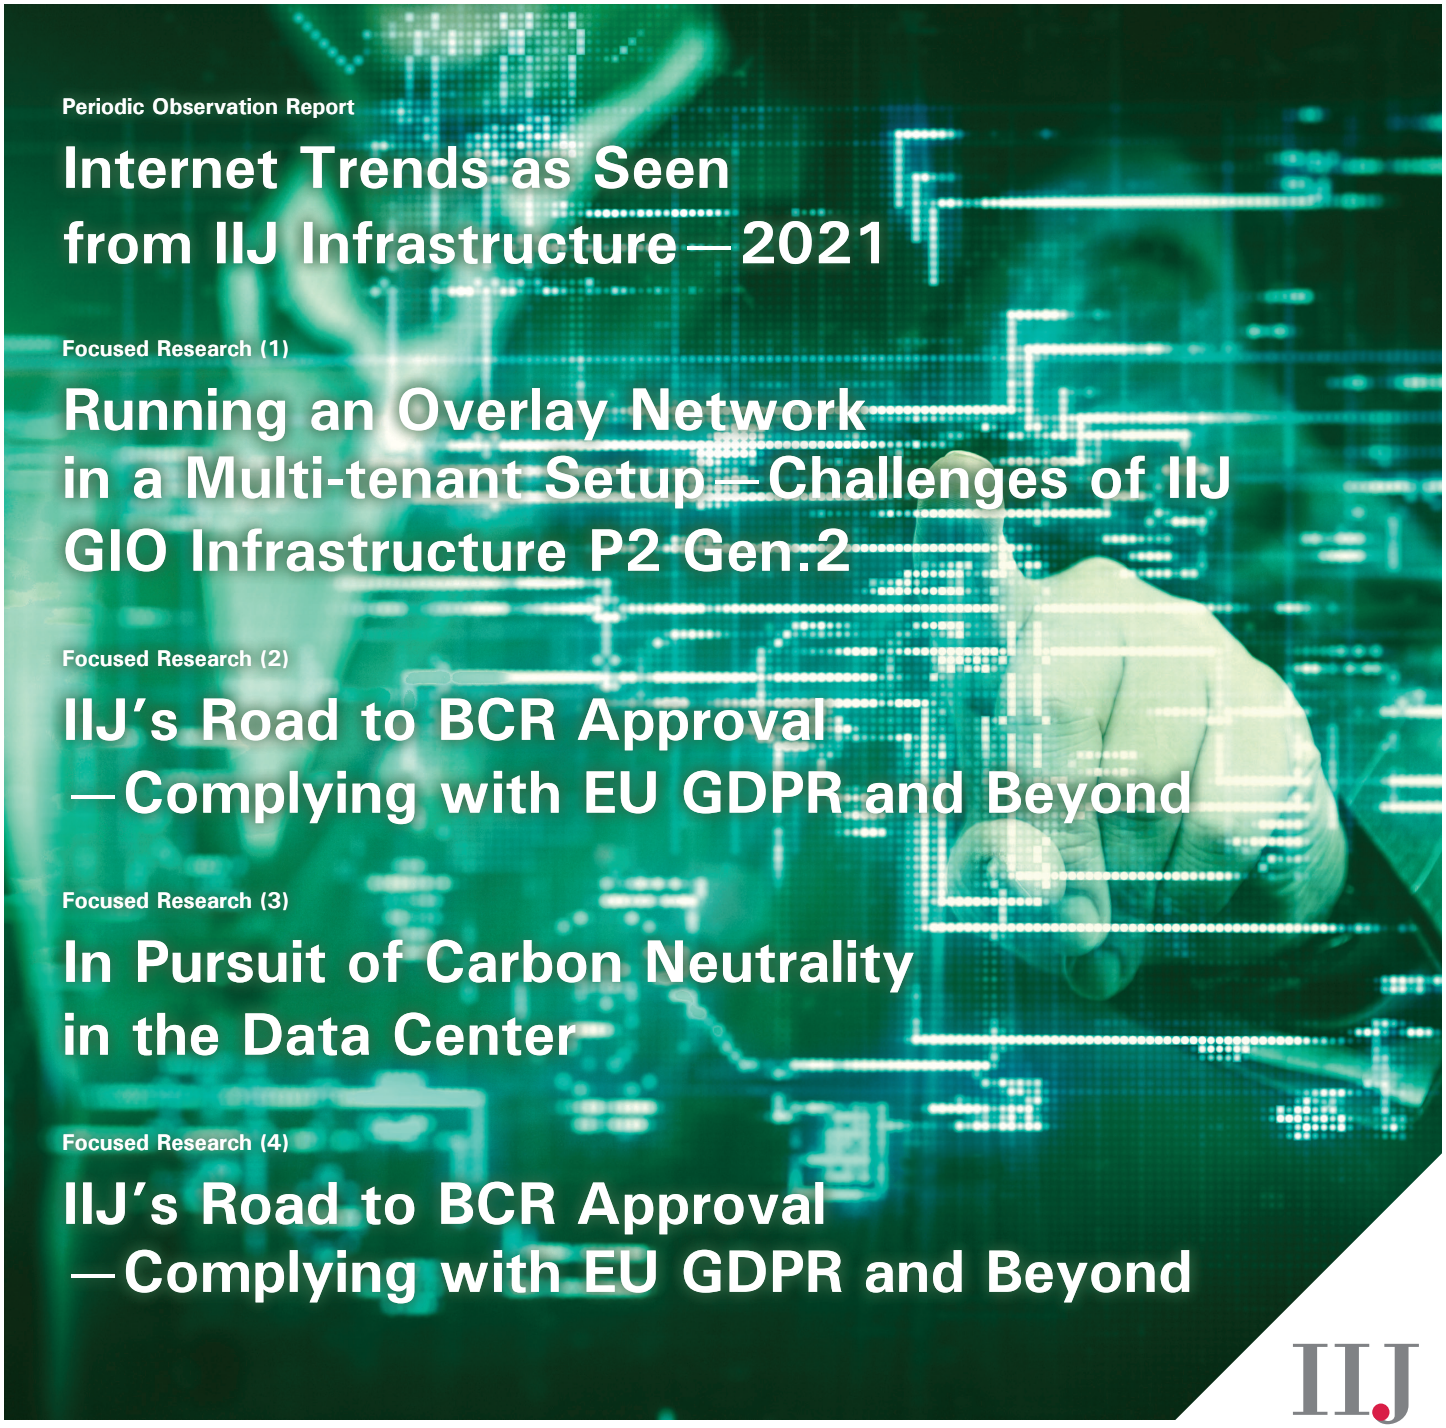

Periodic Observation Report

Internet Trends as Seen from IIJ Infrastructure—2021

Focused Research (1)

Running an Overlay Network in a Multi-tenant Setup—Challenges of IIJ GIO Infrastructure P2 Gen.2

Focused Research (2)

IIJ's Road to BCR Approval —Complying with EU GDPR and Beyond

Focused Research (3)

In Pursuit of Carbon Neutrality in the Data Center

Focused Research (4)

IIJ's Road to BCR Approval —Complying with EU GDPR and Beyond

IIJ
Internet Initiative Japan

Internet Infrastructure Review

February 2022 Vol.53

Executive Summary

To say that modern society runs on information and communications technology (ICT) is no exaggeration. Computers and communication networks are no longer absent from any of our social activities, and we reap the benefits of ICT in terms of increased sophistication and efficiency. What constitutes the infrastructure that supports the workings of society? Energy, transport, government services, finance—the list goes on. Yet none of these would function without ICT. Against this backdrop, Japan is no different from the rest of the world in experiencing the impact on financial and communications services when faults cause system outages at major financial institutions and communications carriers. In our highly information-oriented society, there is an increasing need for reliability with respect to information and communications, and the demands of governance with respect to the businesses that ensure that reliability are also growing. There is no room for doubt about the Internet's role as crucial social infrastructure, and we are well aware that, as part of this, IIJ's networks also play a role underpinning society. We hope to continue meeting society's expectations, and to that end we strive to develop technologies to ensure that we can continue to provide highly reliable services.

The IIR introduces the wide range of technology that IIJ researches and develops, comprising periodic observation reports that provide an outline of various data IIJ obtains through the daily operation of services, as well as focused research examining specific areas of technology.

Our periodic observation report in Chapter 1 presents the 2021 edition of our look at Internet trends as seen from IIJ's infrastructure. The report covers IPv4 and IPv6 routes on the Internet, an analysis of DNS queries obtained from the full resolver provided by IIJ, IPv6 and mobile traffic, and an analysis of traffic during the Tokyo Olympics. As the Internet expands, all sorts of resources and traffic continue to grow, and we are also observing a steady shift in protocols—a rise in AAAA record and HTTPS record DNS queries, for instance, and a rise in absolute IPv6 traffic levels.

The focused research report in Chapter 2 discusses the challenges encountered with the network newly developed for IIJ GIO Infrastructure P2 Gen.2, IIJ's new cloud service released in October 2021, as well as the outlook ahead. While using VMware NSX-T, the engineers also worked in-house to create mechanisms to efficiently operate the infrastructure, which are used for monitoring and capacity planning. The technology is currently used to build networks within data centers and between data centers, but development is ongoing with the aim of extending its applications to connections with distributed edge-computing resources.

Our second focused research report in Chapter 3, titled "In Pursuit of Carbon Neutrality in the Data Center," describes our initiatives at IIJ's Matsue Data Center Park and Shiroi Data Center Campus, offering insights about technologies IIJ has actually deployed at the facilities. Matsue is host to Japan's first commercially operating outside-air cooled modular data center, which runs on a three-phase four-wire power supply, while Shiroi features direct outside-air cooling, system modules, AI control, lithium-ion storage batteries, and more. With carbon neutrality now a major topic, I think you will find the report highly interesting.

Chapter 4 presents our third focused research report on the IIJ's Group's BCRs (Binding Corporate Rules). The IIJ Group created its BCRs in an effort to comply with the EU's GDPR (General Data Protection Regulation), and on August 5, 2021, these BCRs were approved by the competent data protection authority in Germany. In addition to BCRs and the GDPR, the report looks at global efforts around personal data protection, and walks the reader through the IIJ's BCR approval process. The story is a valuable record of events. After IIJ submitted the BCRs for approval in the UK in 2016, the UK withdrew from the EU, and the approval eventually ended up coming from Germany.

Through activities such as these, IIJ strives to improve and develop its services on a daily basis while maintaining the stability of the Internet. We will continue to provide a variety of services and solutions that our customers can take full advantage of as infrastructure for their corporate activities.

Junichi Shimagami

Mr. Shimagami is a Senior Executive Officer and the CTO of IIJ. His interest in the Internet led to him joining IIJ in September 1996. After engaging in the design and construction of the A-Bone Asia region network spearheaded by IIJ, as well as IIJ's backbone network, he was put in charge of IIJ network services. Since 2015, he has been responsible for network, cloud, and security technology across the board as CTO. In April 2017, he became chairman of the Telecom Services Association of Japan MVNO Council.

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.

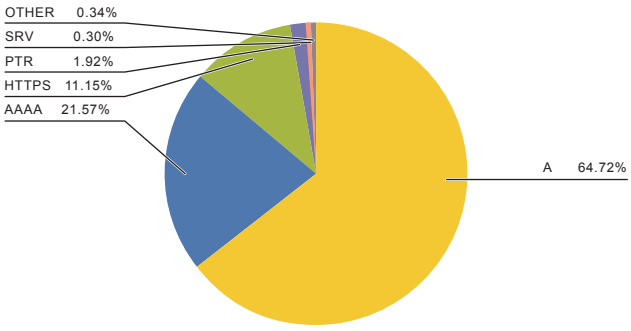

Figure 1: IPv4-based Queries from Clients

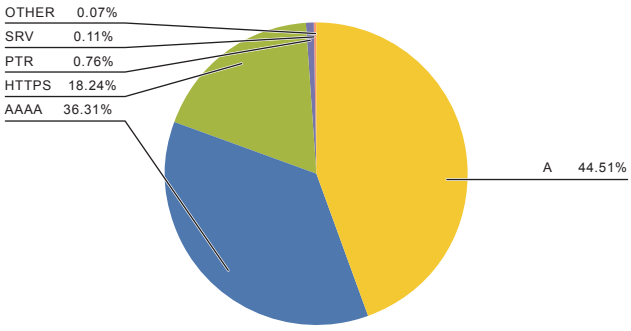

Figure 2: IPv6-based Queries from Clients

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.

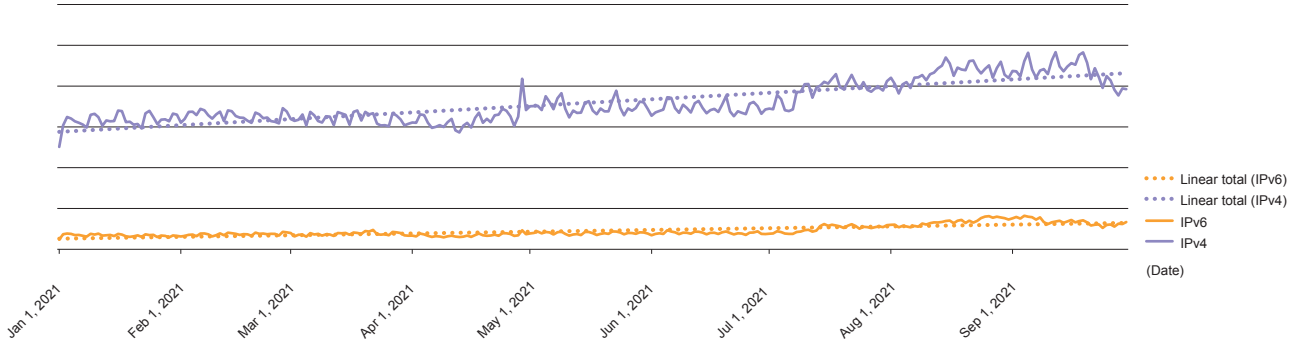

Figure 3: Traffic Measured on Backbone Routes at IIJ's Core POPs

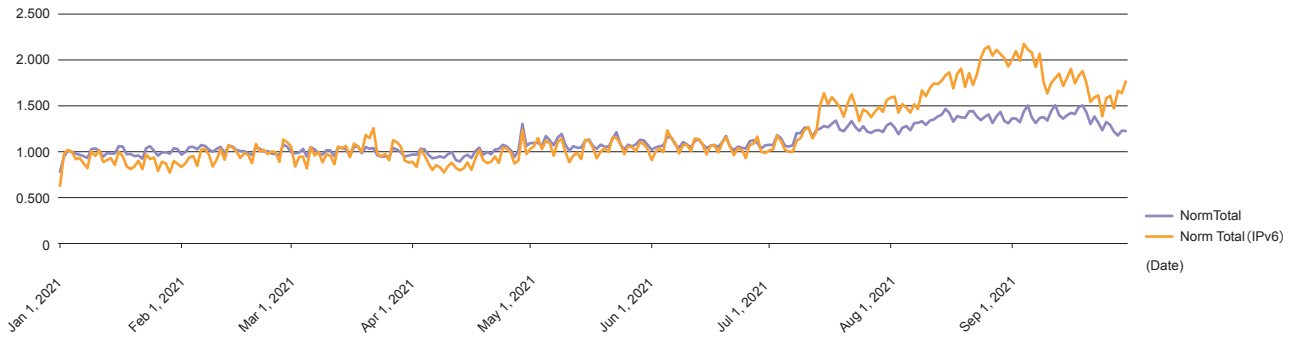

Figure 4: Traffic Indexed to 1 as of January 4

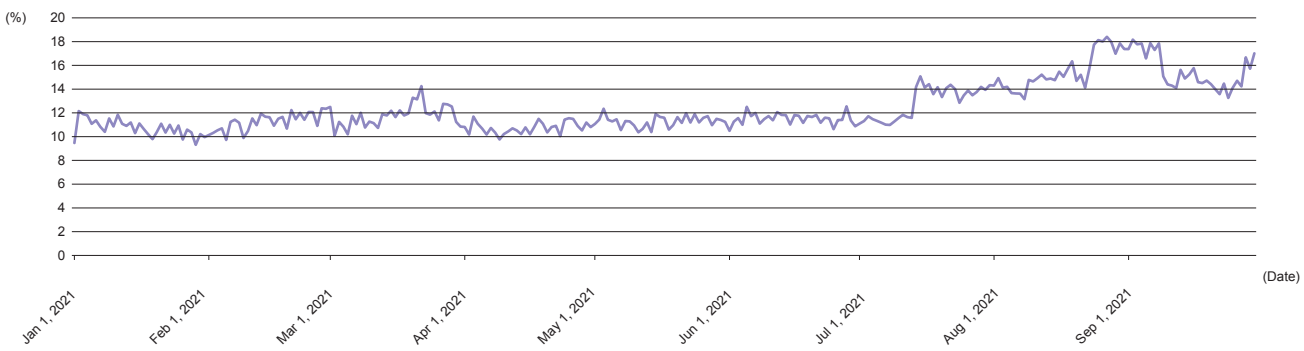

Figure 5: IPv6 as a Proportion of Total Traffic

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%

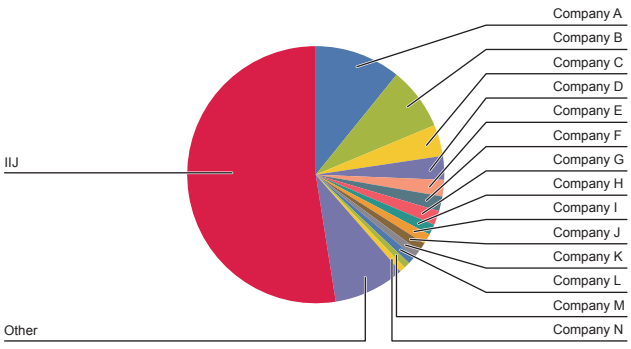

Figure 6: Annual Average IPv6 Traffic by Source Organization (BGP AS Number)

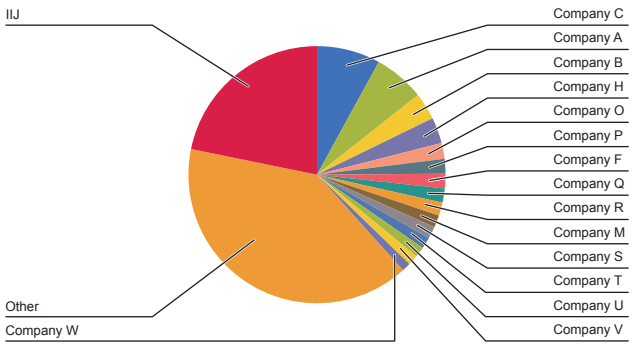

Figure 7: Annual Average IPv4 Traffic by Source Organization (BGP AS Number)

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.

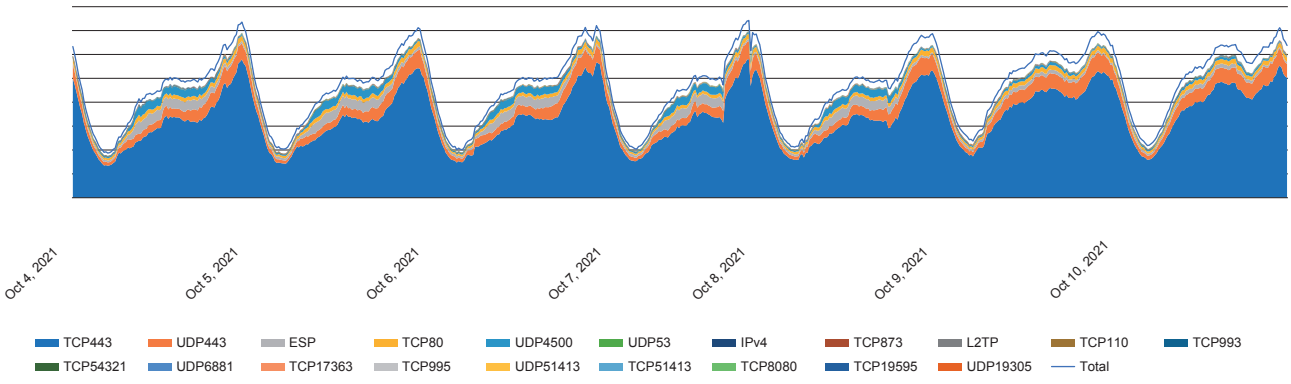

Figure 8: Breakdown of IPv6 Traffic by Source Port Number

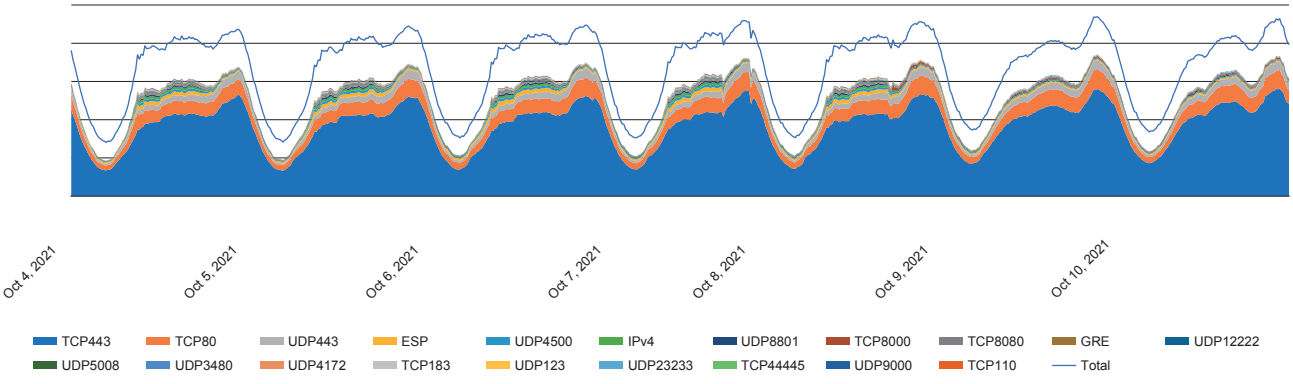

Figure 9: Breakdown of IPv4 Traffic by Source Port Number

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.

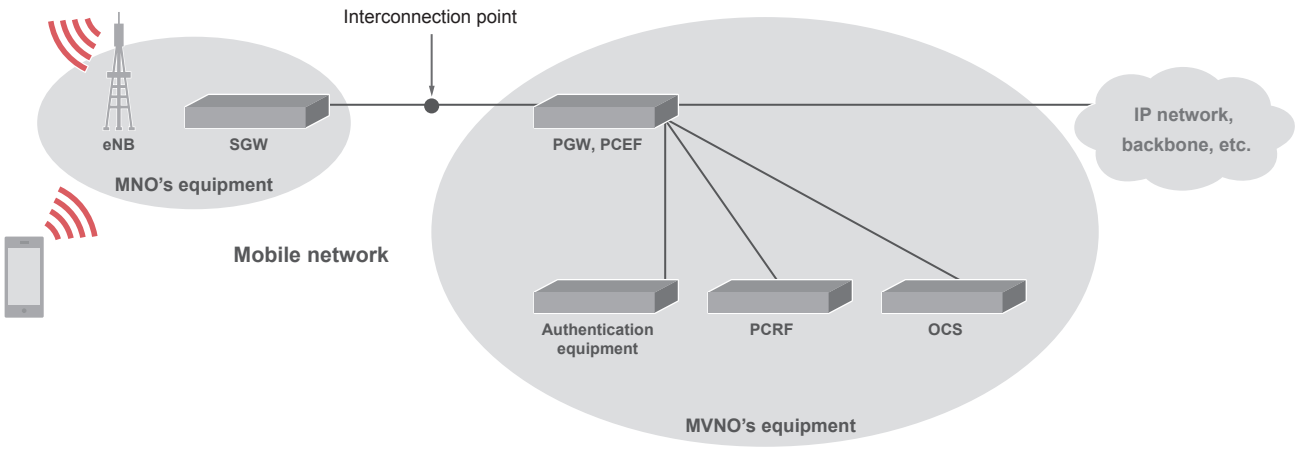

Figure 10: Connection Sequence from Mobile Phone to MVNO Network

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.

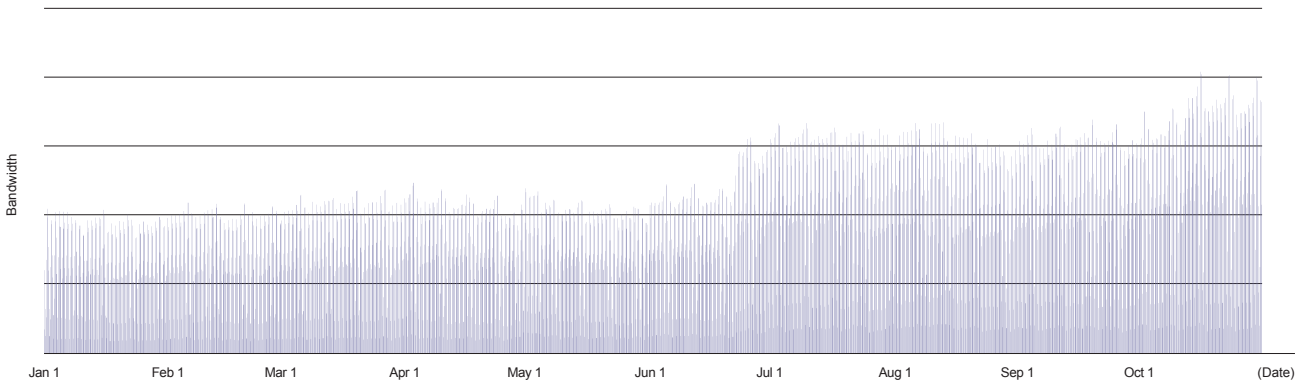

Figure 11:IIJmio Internet-bound Traffic Volume (Jan–Oct 2021)

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.

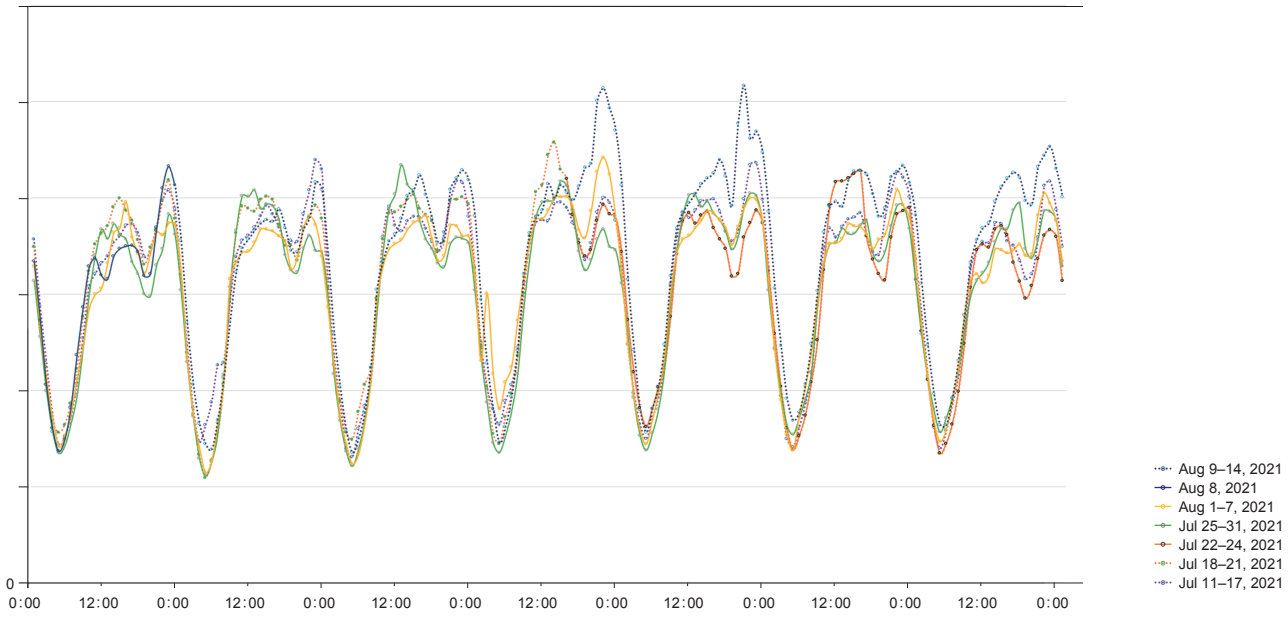

Figure 12: Tokyo/Osaka Traffic

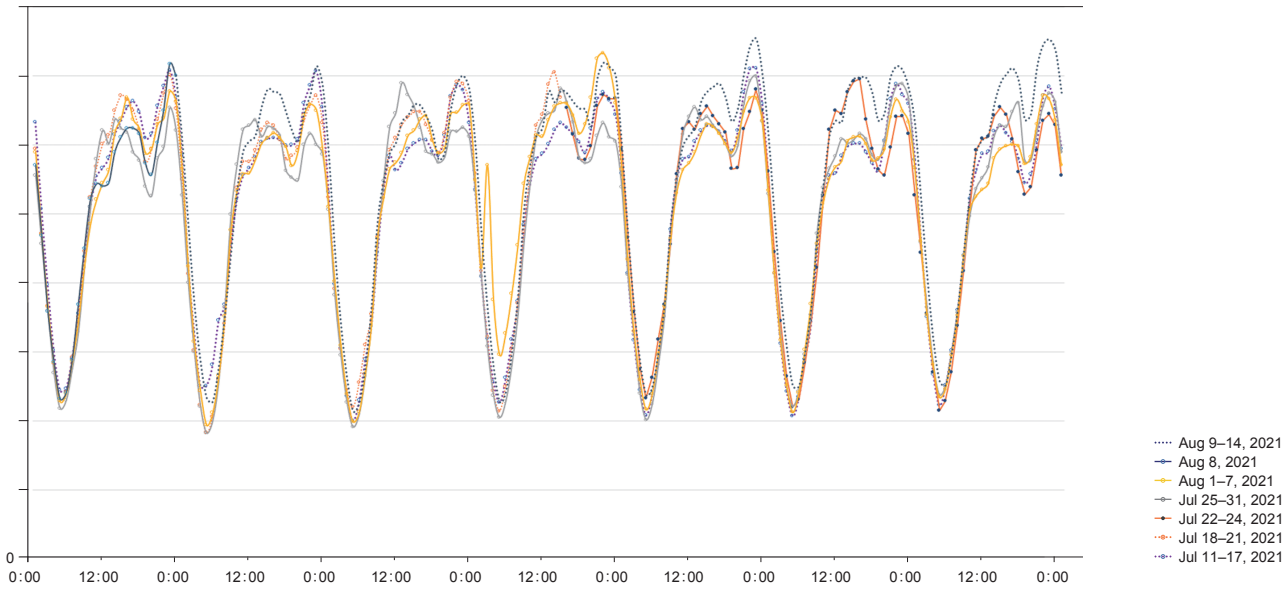

Figure 13: Domestic Interconnection Traffic

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

Running an Overlay Network in a Multi-tenant Setup —Challenges of IIJ GIO Infrastructure P2 Gen.2

2.1 P2 Gen.2—IIJ's new-generation IaaS

On October 1, 2021, IIJ launched IIJ GIO Infrastructure P2 Gen.2 (P2 Gen.2 for short), a new IaaS that fully combines and takes advantage of the characteristics of the public IaaS and private IaaS developed and provided under the IIJ GIO brand. P2 Gen.2 is a new-generation IaaS designed to confer the full benefits of public clouds while also taking advantage of VMware vSphere-based private clouds.

To ensure customers can continue using P2 Gen.2 into the future, we adopted an architecture geared toward the next generation of services. This report looks at the development of the P2 Gen.2 service and the ideals and reality of its operation in the field with a focus on the network.

Our aim in developing and running P2 Gen.2 is to provide a service platform that will continue to be useful a decade or more from now. Today's rapidly changing world makes it impossible to predict what will happen 10 years ahead. Yet we believe that services should remain available for as long as customers want to use them. The development of new systems based on cloud technology under the banner of DX (digital transformation) is on the rise. But maintaining and modifying existing systems to support DX is also a key mission for enterprise IT teams. We developed P2 Gen.2 to address this challenge by providing a third option for cloud computing that can not only accommodate on-premise systems and private clouds that remain in operation but also make it easy to migrate from public clouds as well as.

P2 Gen.2 is an IaaS offering that allows customers to use resources in the form of flexible server resources. The resource pool is virtualized, customers focus on system operations within the resource pool, and IIJ runs all of the functions in the layers below it. With IaaS run in individual VM instances, the size of the VM is made to conform to the instance model specified by the cloud service provider, but with P2 Gen.2, customers are free to create VMs in the resource pool. This means that customers can transfer

the specification of the machines running in their own environments as is. Naturally, this design makes it possible to bring in images or migrate VMs from the existing environment as is via V2V and P2V migrations.

P2 Gen.2 provides an image provision mechanism for deploying VMs; backup, network, and migration functions; and file servers for storing files. The hypervisor and hardware such as servers, storage, and networks are abstracted in a form that customers are able to visualize, so they do not need to worry about these elements. This eliminates the workload encountered when renewing the hypervisor and hardware that results from handing ESXi directly over to the customer, something that had been a major issue with the Virtualization Platform VW Series, a vSphere-based service provided on IIJ GIO.

We adopted VMware Cloud Director (vCD), a VMware product for service providers, and hid the hypervisor (vSphere) layer from users. This made it possible to offer resource control permissions that are as flexible as with vSphere while allowing IIJ, as service provider, to take on the hypervisor and hardware lifecycle management role. Defining a new joint-responsibility model like this made it possible for IIJ to manage and operate the hypervisor network. We adopted VMware NSX-T Data Center (NSX-T), which enables integration with vSphere, to run the hypervisor layer network efficiently, and this significantly improved the IaaS network. With P2 Gen.2, in Layer 3 an NSX-T overlay network sits on top of an IP fabric underlay network, and each tenant's network is completely separated from the others and provided as a VPC (Virtual Private Cloud). Combining this with the operational knowledge in large-scale server pools we amassed through IIJ GIO, we have made it possible to allocate resources to users in a manner independent of the allocation of physical computing resources (CPU, memory, storage) (Figure 1).

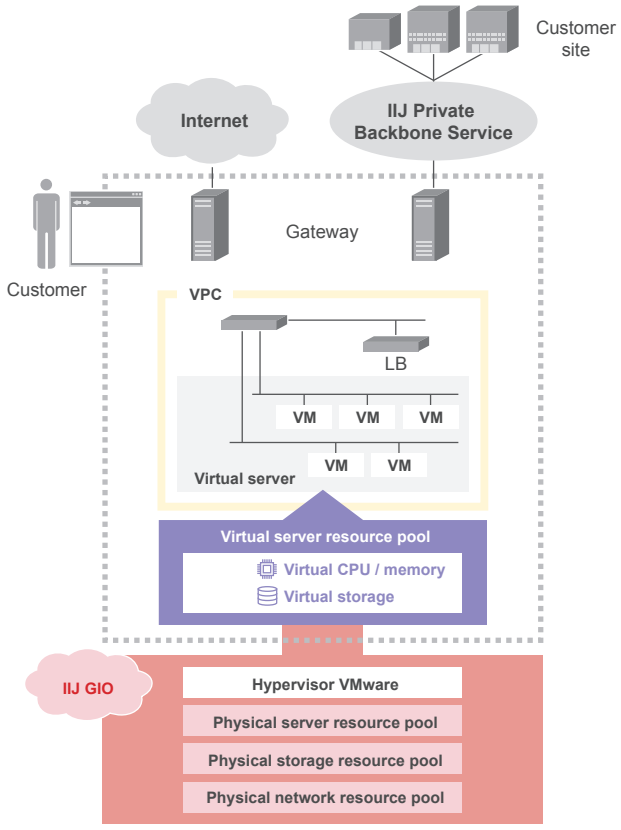

Figure 1: Flexible Server Resources Overview

2.2 Overlay Networks Using SDN Technology

Overlay networks are nothing new. Along with VLAN (Virtual LAN), they are a long-standing network virtualization method. In the past, they were often used for point-to-point connections due to protocol specifications/constraints and for operational considerations, but with the evolution of SDN (software-defined network) technology, they can now be used in large-scale networks. Encapsulation protocols like L2TPv3, VXLAN, and Geneve are used to build virtual L2 networks. SDN technology uses software to control the networks, and is made up of the control plane, which controls what happens between devices, and the data plane, which forwards the actual frames and packets. Using software to control what happens between devices makes it possible to manage the configuration centrally. This creates an environment in which administrators can automatically configure devices according to how the network is defined, instead of having to change the configuration of each individual device, so configurations can be changed quickly to meet the requirements of large-scale networks. This makes it possible to have more than 4,096 network segments, the limit imposed by IEEE 802.1Q VLAN, something that has posed issues in large-scale networks that use many segments. It also enables the efficient control of north-south traffic (traffic between devices and the network center) and east-west traffic (traffic among devices) within virtualized networks (Figure 2).

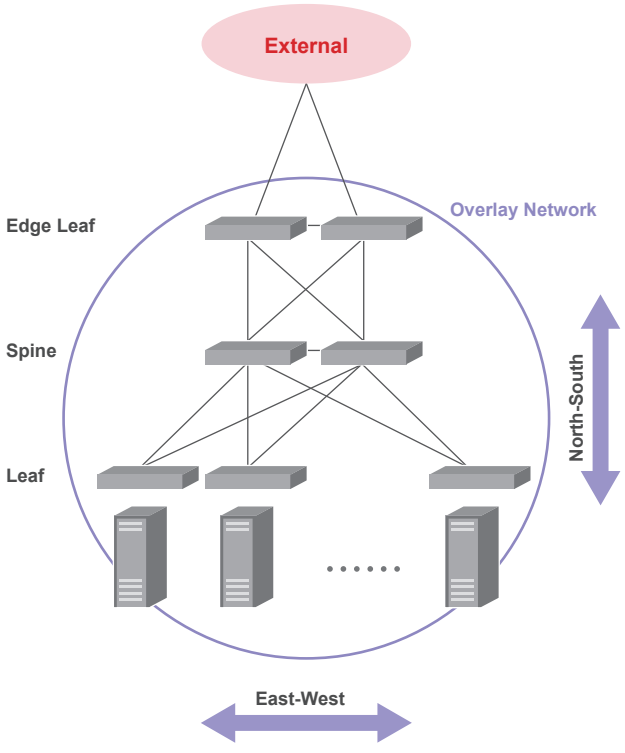

Figure 2: Overlay Network in a Large-scale Environment

2.3 Benefits of Overlay Networks with VMware NSX-T

Along with server virtualization technology, network virtualization with overlay networks is one of the core technologies underpinning the P2 Gen.2 infrastructure. Service networks can be configured separately from the underlay network that lives on the hardware layer. Using NSX-T for network virtualization on the IaaS layer of P2 Gen.2 provides a flexible network configuration and loose coupling with the physical layer. When setting up networks with a view to expansion from minimum scale up to hyperscale, they tend to become complicated in terms of physical device specs because design and sizing need to be chosen in anticipation of that expansion. But being able to control expansion via overlays confers considerable benefits on the operational front.

An IaaS also needs to handle multi-site configurations across multiple data centers. Using infrastructure across multiple, physically separated sites facilitates disaster preparedness and makes it possible to obtain the necessary computing resources without needing a large-scale facility. Virtualizing the network between sites and making it an overlay network enables you to configure networks independent of their location. Physically separated environments involve more considerations to take into account than with single-site configurations, such as network latency and bandwidth between sites, and this increases design and operational difficulty, but ensuring proper resource monitoring and control of systems based on the monitoring data makes it possible to set up flexible networks.

Network virtualization remains uncommon despite these benefits. With traditional networks where the configuration was expected to remain static, administrators were able to continue maintaining and running systems without the benefits of network virtualization. In recent years, however, there is a need for IT infrastructure as a whole to adapt flexibly and swiftly to rapidly changing and unpredictable requirements. Network virtualization will likely become more and more important in this landscape.

IaaS solutions need to be multi-tenant (accommodate multiple customers). Centralized control and automation via software is also needed in order to provide functions separately to the tenants and to maintain security and the specified service level. Network virtualization makes it easier to link together the other components of the IaaS solution as well as monitoring systems, facilitating proper control of a multi-tenant environment in which each tenant operates their system in accord with their system requirements. NSX-T also implements VRF (virtual routing and forwarding), and VRF can be used to enhance equipment efficiency. In multi-tenant environments, the need to divide up routing domains among the tenants posed a resource constraint. VRF makes it possible to increase infrastructure consolidation.

Interface control is centralized into NSX Manager, NSX-T's management component, making API-based operations and automation easy. This also helps to reduce development hours spent on making the service's backend applications work together.

When an overlay network is set up as an IaaS solution, network exits are needed. NSX-T also cannot communicate outside the NSX-T domain (under its control) with the NSX-T encapsulation protocol Geneve, so to communicate with the external world, you need to convert it to a VLAN. We had not fully automated external connections with IIJ GIO, and a major issue was that users were unable to change network configuration in a complete form when migrating systems. P2 Gen.2 improves this situation by automating route configuration for connections with the IIJ Private Backbone Service. NSX-T can also be run in a bare metal configuration with NSX Edge installed directly on a server with no virtualization, and while this is fine in performance terms, it is not really up to the task of multi-tenant configurations. Many challenges also present themselves particularly with external connections, such as downtime during software upgrades. So many necessary areas of improvement remain, and we continue to work on resolving these issues.

2.4 Operational Issues and Solutions

Ahead of the October 1, 2021 release, in 2020 we conducted a small-scale rollout of the NSX-T SDN-based overlay network, along with NSX-T's NFV (network function virtualization) capabilities, for some data center functions. This revealed a whole host of issues, which we continue working to address as part of our operations. Below, we go over the issues we discovered and their solutions.

Creating an overlay network using SDN requires a more complicated configuration than with conventional networks. The configuration has two layers: the physical layer (underlay) and the IaaS layer (overlay, logical layer).

Another point is that to get the most out of the NSX-T architecture, you need to be well-versed in vSphere operations, so in addition to networking, the engineers also need to be able to work with server, hypervisor, and virtualization technologies. This is a departure from the past practice of specializing in specific areas—à la conventional network engineers and server engineers—and instead, the engineers need to be familiar with the entire technology stack used in building the infrastructure. When we started out, our team was small but included people with a strong knowledge of VMware technologies and people with experience running large-scale service infrastructure, so as a team, we had the capabilities to handle ongoing development and operations. The most important part of getting engineers with differing areas of expertise to work both autonomously and in collaboration with each other as a team is to have a vision for your products and services and strive continuously to make them better based on the DevOps philosophy.

With P2 Gen.2, we revisited infrastructure control from scratch to ensure it would work the way we needed it to. We put as much effort as possible into streamlining operations and automation. Under DevOps principles, we built a Git-based CI/CD pipeline, and the engineers control the majority of elements that make up the IaaS,

not just the VMs, under a policy of making changes to the operations and configuration of the infrastructure without directly manipulating the environment (Figure 3).

While there is a lot of knowledge and software out there for controlling VMware products, the de facto standard in server virtualization, it is not quite adequate for total control of multi-tenant infrastructure in the form of a SDDC (Software-Defined Datacenter), so some of the tools and libraries we use were developed from scratch.

We are not only simply carrying out operations tasks through these efforts. We have also laid the groundwork for implementing SRE (site reliability engineering).

Taking on board agile values and methods like this can facilitate more flexible and agile infrastructure operations, which can otherwise tend to be fixed and inflexible, but you also need to be keenly aware of the opposing relationship between agile principles and facility operations. Procuring physical equipment involves delivery lead times, and setting up physical equipment in the data center is not something that you can do via software control alone. Configuring facilities involves some level of physical human interaction, so making sure the associated lead times do not become a bottleneck is key. Having an abundant inventory of resources can solve this issue, but this then leads to inefficiencies because the overall resource utilization rate falls. Also, out-of-resources conditions are unacceptable in cloud services. So resource management, monitoring, and demand forecasting are crucial. Monitoring and demand forecasting should be done not just on the overlay, it should also aggregate and facilitate visualization of the various underlay data sources, including servers, network equipment, and storage. Since the October 1, 2021 release, we have been collecting and analyzing data on actual usage conditions to enable just-in-time resource inventory deployments based on demand forecasting.

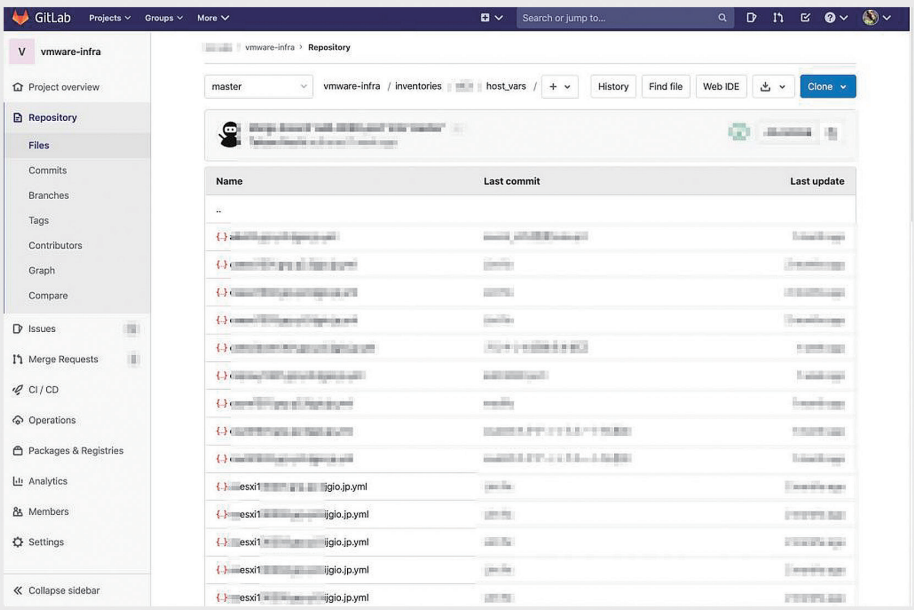

Figure 3: The IaaS Components are Managed via Git

2.5 Looking Ahead

P2 Gen.2 still only has the minimum functionality, and we plan to continue developing it to serve customer needs. The use of computing resources in the form of cloud-native IT systems has become commonplace, and a more agile digital infrastructure is now being demanded. Our first priority is to implement the features needed to ensure that customers will continue to use the service, but we also want to be able to address future needs beyond that.

In addition to migrations into the cloud, the need for edge computing and distributed computing that deploys computing resources close to where users are is also rising. IIJ's role is to connect the cloud at the core, the distributed edge nodes, and the network in between. These interconnected elements will allow IIJ to deliver high-quality services on edge nodes closer to where the customer is actually located. In addition to conventional networks, 5G mobile networks and the like can also be used to connect the resources. Linking the core and the edges requires not only VMware-based workflow management but also application container-based resource management, so also providing resources in the form of Containers as a Service (CaaS), allowing applications to be run inside containers, might be the optimal solution for customers (Figure 4).

Software control that brings the network closer to the status of an application is needed here. This goes beyond just resource services that provide VMs, and I think it can be achieved by evolving the network from the current

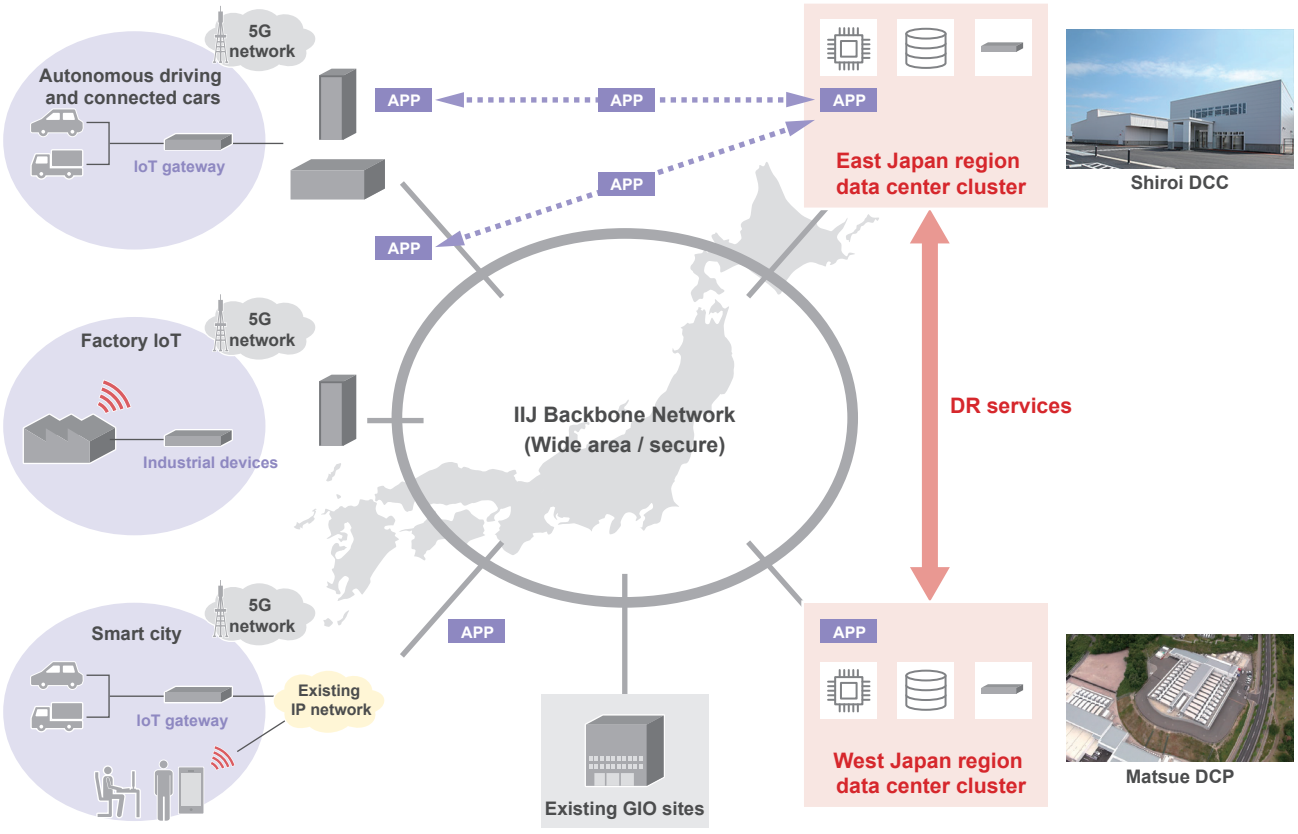

Figure 4: Looking Ahead

architecture into something that is even less dependent on physical resources. This will be difficult to accomplish simply by deploying resources in a compact form that can also be deployed on the edges. You also need to be able to configure the network flexibly so that it works together with the other resources.

Personally, I am interested in SmartNICs as a technology suitable for this sort of edge computing and distributed computing architecture. VMware is working on Project Monterey, which runs ESXi on Arm-processor based SmartNICs. The hope is that managing SmartNICs as NSX-T domains will make it possible to control and deploy networks in a manner that makes them equivalent to applications.

To make this sort of world a reality, we are also working on a project to use the technology stack we have developed not just for data center purposes but also as an edge computing platform. We have set up a micro data center outdoors at IIJ's Shiroi Data Center Campus, and we are conducting tests in the area of cloud and edge computing as well as multi-access edge computing (MEC), with our objective being to develop the technology further along with local 5G, IoT, and other use cases. My hope is that evolving network services in this way will make it possible to provide customers with unique infrastructure of the sort that only IIJ can deliver.

Takehiro Yamamoto
Infrastructure Service Design Manager, Cloud Services Division 1, IIJ System Cloud Division.
Mr. Yamamoto is engaged in the planning, development, and operation of cloud-related services.

In Pursuit of Carbon Neutrality in the Data Center

3.1 Introduction

Over 120 countries have declared their intention to achieve carbon neutrality—achieve zero greenhouse gas (CO2, methane, N2O, chlorofluorocarbons [CFCs]) emissions—by 2050, a goal of the international climate change treaty known as the Paris Agreement. The Japanese government also declared its 2050 carbon neutrality goal in October 2020. And in the Green Growth Strategy it unveiled in December 2020, it set goals for a wide range of industries, with the following goals being laid out for data centers.

- By 2030, all newly constructed data centers should consumer 30% less energy and be partly powered by renewable energy
- Aim for carbon-neutral data centers by 2040

Data center carbon neutrality is to be achieved primarily by reducing CO2 emissions from power consumption to zero. This will entail costs including for the purchase of renewable electricity, technology development, and capital investment, but achieving carbon neutrality early will help companies differentiate from peers and enhance service value. Carbon neutrality is also starting to appear among procurement criteria for products and services, and equity markets are calling on companies to disclose information on how they are effectively impacting climate change rather than simply mechanically disclosing data, and these developments also no doubt create an incentive for carbon neutrality.

Data center carbon neutrality will be achieved via two avenues: saving energy by using it more efficiently and using renewable energy not tied to CO2 emissions. Matsue Data Center Park (Matsue DCP) and Shiroi Data Center Campus (Shiroi DCC), which IIJ built and operates, feature highly energy-efficient equipment and outside-air cooling systems the technology for which was developed and tested with high energy-saving goals in mind, and the resulting energy savings are reducing IIJ's greenhouse gas emissions.

In addition to such energy savings, we also plan to make increasing use of renewable energy ahead by setting up carbon-neutral data centers so that renewable energy generation systems and the data center can work in concert with each other.

Below, we go over our energy-saving initiatives at Matsue DCP and Shiroi DCC, where we have achieved strong energy-saving performance, and we round out the discussion by presenting a reference model for carbon neutrality in the data center.

3.2 Outcomes at Matsue Data Center Park

On April 26, 2011, we opened Japan's first commercial modular data center with outside-air cooling systems, Matsue Data Center Park in Matsue-shi, Shimane Prefecture. Matsue DCP features IZmo units, IT modules developed by IIJ and imbued with its extensive data center operational knowhow. To expand into more enterprises in Japan and abroad and tap into new markets, we also operate co-IZmo/I modular data centers, which use indirect outside-air cooling systems. Matsue DCP has been in stable operation since being opened as the core facility for IIJ's GIO cloud services, and with 90% of the server-housing containers now installed, the total number of servers running is in the tens of thousands. Site 2, added next to Site 1 in 2013, introduced co-IZmo/I units in addition to IZmo and uses a three-phase four-wire configuration to power its servers and other IT equipment, reducing power distribution losses and thus saving energy (Figure 1). Data centers generally house large-capacity electrical equipment and air conditioning systems to create an environment in which large numbers of servers and other IT equipment can be installed efficiently. IT equipment is the biggest consumer of data center power, but air conditioning systems are right behind it. So we needed to rethink the use of conventional air conditioning systems and introduce new systems that consume less power. Based on past demonstration testing, IIJ determined that outside-air cooling systems that use outside air to reduce power consumption and do not require cooling towers and the like would be suited to its next-generation data centers. Outside-air cooling, however, involves the exchange of a large volume of air, so the intakes and outlets need to be installed directly on the server rooms. This raises a number of difficult problems to do with building structure when seeking to install such systems on existing buildings. So we developed the IZmo IT module, in which the air ducts are integrated with the server room.

PUE (Power Usage Effectiveness), a metric developed by the Green Grid, is often used as a measure of how effectively a data center uses power. Figure 2 shows the PUE equation.

Theoretically, the smallest (best) PUE score is 1.0, which occurs when zero power is consumed by air conditioning equipment etc. Figure 3 shows five years of measurements for the IT-module-only metric pPUE (partial PUE, ignores energy losses from shared resources etc.).

The summer pPUE readings are higher than for other seasons because the air conditioning modules run in a

Figure 1: Matsue Data Center Park Site 1 and Site 2
(Photo of Matsue Data Center Park taken around 2019)

$$PUE = \frac{\text{Total facility energy (energy used by IT equipment + energy used by air conditioners etc.)}}{\text{Energy used by IT equipment}}$$

Figure 2: The PUE Equation

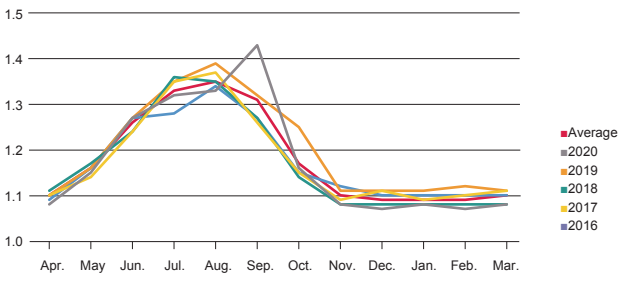

Figure 3: Annual pPUE Scores for the Past Five Years

high-power-consumption mode that does not use outside air. In comparison, spring and autumn pPUE readings are around 1.1 because the modules run in a low-power mode that uses outside air. In winter, the modules mix IT equipment exhaust and outside air to achieve the right temperature, and since this mode also consumes less power, pPUE is around 1.1. The average over the past five years is 1.18.

IZmo and co-IZmo/I

The IZmo IT module is a non-building container conforming to technical advice issued by Japan's Ministry of Land, Infrastructure, Transport and Tourism in March 2011. The modules are easy to set up because there is no need to apply for building permission. With conventional data centers, the servers had to be unpacked one by one, installed in a rack, and wired up. But with containerized data centers, the servers are installed into IT modules at the server factory, and the IT modules are transported by truck for installation at the site. This eliminates box and packaging waste, makes better use of resources and is better for the environment, and helps reduce CO2 emissions from equipment transportation. Recently, when replacing IT equipment, instead of removing the entire container, operators can opt to remove all of the server racks from the container and take them to the server factory, where the new equipment is installed in the racks so they can then be reinstalled in the containers.

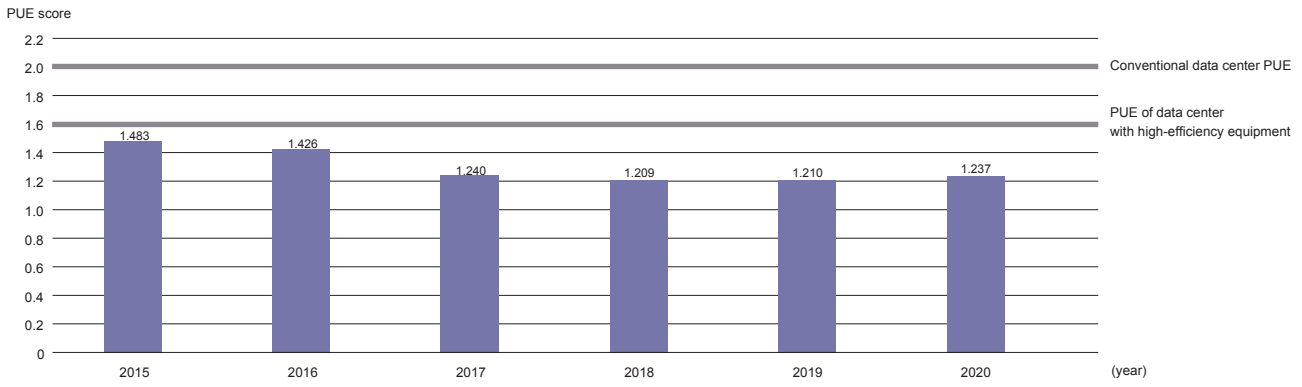

Figure 4: Annual Average PUE for Matsue DCP as a Whole

Figure 5: IZmo and co-IZmo/I

IZmo uses direct outside-air cooling and automatically switches between three modes depending on outside air temperature and humidity. It controls the temperature at the intake port within the IT module to meet the recommended temperature and humidity conditions set by Ashrae (American Society of Heating, Refrigerating and Air Conditioning Engineers) in 2008.

co-IZmo/I uses indirect outside-air cooling. In contrast with direct outside-air cooling, outside air is not channeled directly into the container, and instead, internal heat is expelled indirectly via a heat exchanger, so this method can be used even when outside air conditions are poor. We have also exported some of these modules internationally. See Figure 5.

Theoretically, IZmo, which uses direct outside-air cooling, is more efficient (lower pPUE) than co-IZmo/I with its indirect outside-air cooling (pPUE value is low), but as Figure 6 shows, the fine-tuning of control mechanisms and the like make it possible to run co-IZmo/I modules at pPUE scores similar to those of IZmo.

Three-Phase Four-Wire System

At Site 2, we introduced a three-phase four-wire uninterruptible power supply (UPS) system, a first among Japanese data centers (Figure 7). Three-phase four-wire power transmission involves three wires for three-phase AC plus a grounded neutral wire. It is more efficient than single-phase AC, uses a lower current, and conductor (wire) size can be reduced by transmitting power at high voltage (low current). It is widely used in Japanese factories. One of the three 400V wires from the UPS can be isolated along with the neutral wire to pull out a 230V supply without a transformer. This is common in data centers overseas where 100V is not required, but the common need for 100V in Japan means that, until now, transformers have commonly been installed in three-phase three-wire UPSs to transform power from 400V to 100V. Since servers can operate at any voltage from 100V to 230V, and because IIJ uses a lot of servers for its cloud infrastructure, we chose a three-phase four-wire system, eliminating the need for transformers.

Given that power (W) = current (A) x voltage (V), a reduction in wire size means that, for any given load capacity (W), the higher the voltage (V), the more you can reduce the current (A). Specified currents for electrical wire (the maximum current it can carry) are set according to the type of insulating sheath and the size of the wire (mm2). So at lower specified currents, wires with a lower permissible current can be used, and the lower a wire's permissible current, the smaller its diameter, which can help reduce the investment cost. Further, voltage drop can be reduced because three-phase four-wire systems are even less susceptible to voltage drop than three-phase three-wire systems, in which the voltage drop decreases as the current decreases. At Matsue DCP, this allows for a lower transmission current between the UPS output and the server input, and it reduces losses because no transformers are used. This results in a theoretical reduction in losses of around 25%.

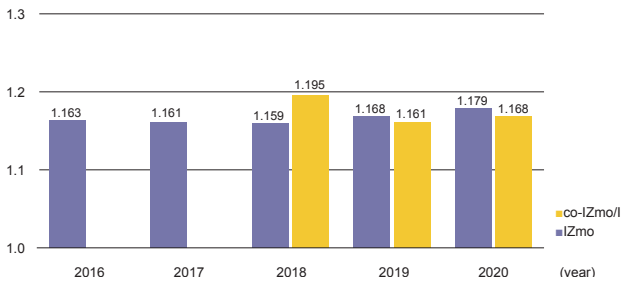

Figure 6: Comparison of IZmo and co-IZmo/I pPUE Scores

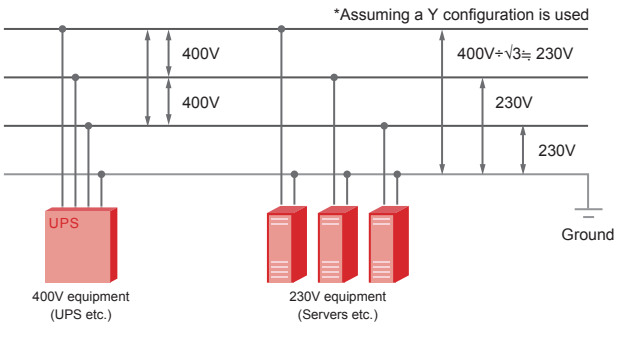

Figure 7: Y-connection Overview of Three-phase Four-wire System

Environmental Initiatives

In terms of environmental initiatives, we have ISO 14001 (EMS, Environmental Management Systems) certification and work to comply with energy conservation laws and regulations.

Matsue DCP operates in accord with its ISO 14001 certification, acquired in 2013. We set annual targets for environmental initiatives. Our past efforts have included going paperless, reviewing the air conditioner temperature settings within buildings and operations rooms, and reducing pPUE by 0.01 per year. Using power more efficiently is one thing, but the EMS certification also requires us to address the risk of the environmental impact on lives and ecosystems from atmospheric pollution caused by fires and the impact of soil pollution caused by the leakage of fuel from emergency power generation facilities. We have therefore prepared response manuals for emergencies such as fires, earthquakes, and equipment accidents, and we regularly conduct response training sessions, environmental management education and training sessions, and environmental compliance assessments. Matsue DCP is in an area that experiences relatively few natural disasters and earthquakes, but we strive to raise environmental awareness through regular education and training.

On the topic of environment, energy, and CO2 emissions, we have compiled data on our energy usage into periodic reports in accord with energy conservation laws and regulations since Site 1 was opened in 2011. Because Matsue DCP's energy usage exceeded 1,500kL (crude oil equivalent) in fiscal 2013, we have also been working on energy saving measures in accord with the facility's classification as a Type 2 Designated Energy Management Factory. We have appointed an Energy Manager and formulated a medium-term plan, and we implement and evaluate energy-reducing initiatives based on the facility's energy usage every year. As the facility has many air conditioning systems, we also comply with Japan's CFC emissions law by measuring the amount of CFC gas leaked and identifying leaks and their causes through inspections and such.

As Matsue DCP's electricity usage increases by the year, efforts to save more energy and further reduce CO2 emissions are becoming increasingly important. We will continue to focus on environmental initiatives including those to reduce energy consumption, use renewable energy, and install renewable-energy equipment.

3.3 Initiatives at Shiroi Data Center Campus

We opened Shiroi Data Center Campus on May 1, 2019 in Shiroi-shi, Chiba Prefecture (Figure 8). Shiroi DCC is a large-scale data center designed to cope with the explosive growth in data center demand for 5G, IoT, AI, cloud services, and the like. It brings together the data center technologies and knowhow that IIJ has developed or evaluated to address issues with data center operations in the past. It is also a system module-type data center where we can actively roll out new technologies.

Shiroi DCC draws on the successes of Matsue DCP and is designed to overcome the obstacles and challenges encountered. On the energy front, it inherits the concept of modular server rooms, direct outside-air cooling systems, three-phase four-wire UPS and bus duct power transmission technology. The entire Shiroi DCC facility runs at an annual average PUE of under 1.2, and we are endeavoring to expand the usage scenarios for its power systems.

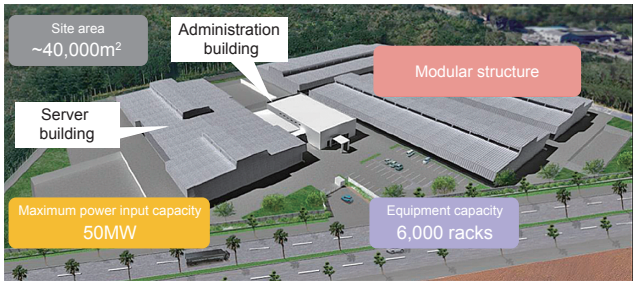

Figure 8: Artist's Impression of Shiroi Data Center Campus Once Completed

Direct Outside-Air Cooling Systems and System Modules

As a new modular data center, Shiroi DCC takes advantage of operational experience from Matsue DCP and knowhow gleaned from containerized data center development, applying this on a floor scale (several hundred racks or so) and using system modules capable of serving needs on a larger scale. The buildings follow a systematized construction scheme based on standard shapes and arrangements for the steel building frames and outer walls, facilitating rapid construction at low cost and high quality (Shiroi DCC phase-1 building: from start of construction, it took four months to have the roof frame in place, eight months to fully complete). All of the internal elements of the system modules—electrical system, air conditioning system, and server racks—are modularized, making it possible to add or replace (upgrade) equipment for each server room individually.

These system modules are also integral to the reduced energy usage of the air conditioning systems at Shiroi DCC, which has a huge number of racks. Using the system modules' characteristically large space, large apertures, and wide span, we created a side-flow direct outside-air cooling system that utilizes the chimney effect.

Figure 9: Test of the Chimney Effect (2012)
*Patent no.: 6153772 (granted June 9, 2017)

Chimney Effect

The direct outdoor-air air conditioning systems in use at Matsue DCP and on IZmo containerized data center units have greatly reduced the annual power consumption attributable to outdoor units, resulting in an annual average PUE of 1.237 for Matsue DCP and annual average pPUE of 1.18 for containers. The systems need ventilation fans to blow cool air into the cold areas year round, however, and the power used by these fans was the next energy-saving issue we faced with the air conditioning systems. To tackle this issue, in 2012 we conducted tests (Figure 9) to see if we could take advantage of the chimney effect (phenomenon whereby the difference in temperature on the inside and outside of a chimney causes an upward current) to cool the servers. If this was possible, we thought, it could greatly reduce the power consumed by the ventilation fans and take us closer to a PUE of 1.0, and if it also made it possible to reduce IT equipment internal fans, this could also help to further reduce overall data center power consumption.

Although the conditions under which the chimneys were installed (weather, air temperature and humidity, wind speed, etc.) did affect the results, our tests showed that airflow volume increased roughly proportionally with increases in the height and load capacity of the chimneys (Figure 10). As a result, we determined that we could assist the air conditioner ventilation fans by grouping several racks into a single module and positioning several modules at the center of the room so that outside air could be drawn in from the sides of the room, and exhaust air from the modules could be expelled from the top of the building. Shiroi DCC is designed with these test results in mind in terms of the system module layout and the direct outside-air cooling systems.

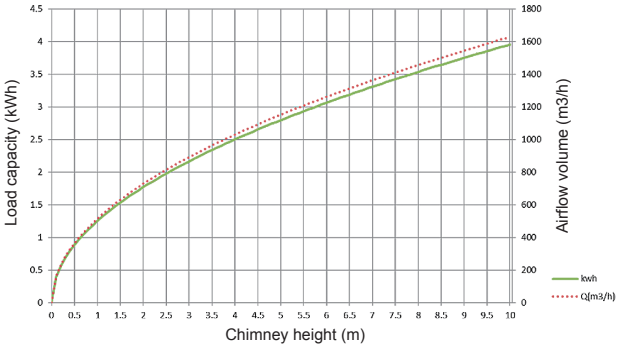

Figure 10: Chimney Effect Test Results
(excerpt: change in airflow volume according to chimney height)

Side-Flow Air Conditioning

Shiroi DCC uses a side-flow system (IDC-SFLOW, a trademark of Takasago Thermal Engineering Co., Ltd.) that supplies cool air directly into the server room via the partition wall between the server room and the air conditioner room (Figure 11).

Data centers housed in buildings often employ floor ventilation with air supplied via the space under a raised, free-access floor. Space is tight, however, and the speed of airflow through the server room inlet (free-access grille) is high. This sort of design also has communications and power cables routed through the same space, which increases air supply resistance and results in a greater loss of pressure. A large pressure loss means more energy is required for air conditioning, so the ventilation fans use more power. Shiroi DCC's wall-based side-flow air conditioning uses the system modules' characteristically large space, large apertures, and wide span as the air supply air route. The inlet openings are designed to be as large as possible so that when air is blown through the server room, air supply resistance is lower and the pressure loss is thus minimized. Shiroi DCC is designed to use an average of 6kW/rack, and the heat generated by the servers can be dissipated by blowing air via the inlets into the server room at a slow airflow rate of 2m/s. This made it possible to greatly reduce (to around a third) the amount of power consumed by the air conditioning system's ventilation fans. The low airflow rate also makes for a better working environment in the server rooms and greatly reduces the stress on workers.

Direct Outside-Air Cooling

Shiroi DCC uses direct outside-air cooling systems that are integrated with the system modules. Air is drawn into the air conditioner room via intake ports on the eaves on the side of the system module, passing through a medium-performance filter, and then mixed with server exhaust heat and supplied into the server room at a set temperature and humidity. And server exhaust heat, in the same amount as the outside air drawn in via the eaves, is expelled via exhaust ports at the top of the building.

Figure 12 shows the basic plan for a large-scale data center created in 2012. At the time, there was skepticism within IIJ about whether we could actually build something like this. But the direct outside-air cooling systems designed on top

Figure 11: Side-flow Air Conditioner Outlets

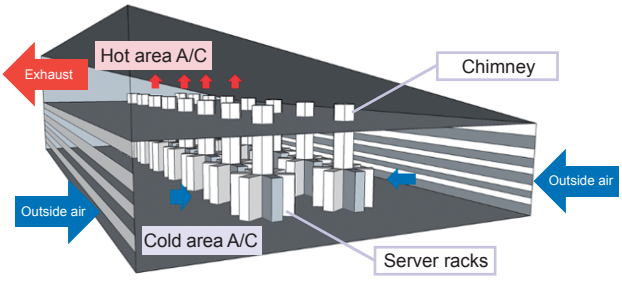

Figure 12: Modular Data Center Design Proposal at Time of Chimney Effect Testing (2012)

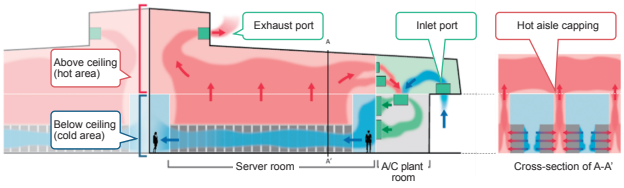

Figure 13: Cross-sectional View of Shiroi DCC System Module (2019)

of knowhow gleaned from Matsue DCP and technical tests on the chimney effect and the like culminated, after almost 10 years of effort, in Shiroi DCC as depicted in Figure 13.

Using AI to Optimize Air Conditioning Systems

At Shiroi DCC, we use AI to find the optimal combination of parameters for not only the air conditioners but the IT equipment as well. We endeavor to optimize air conditioning system operating conditions based on metrics of energy savings, cost savings, and CO2 reductions.

2012 Chillerless Tests and 2013 IT Equipment Adaption Tests

The air conditioning systems on the IZmo units in use at Matsue act as direct outside-air systems during autumn, winter, and spring, but during summer, the outside air is shut off and indoor units are used to cool and dehumidify (circulation mode), so the energy savings in summer are worse than in the other seasons (autumn/winter/spring pPUE = 1.095, summer pPUE = 1.331). In an effort to resolve these issues, we did an assessment of data center health under a scenario of using outside-air cooling year round and not running the external air conditioners, chillers, etc. When Matsue Site 2 went into service in 2013, we developed co-IZmo/D, a new type of containerized data center module that uses year-round outside-air cooling, and started conducting a whole series of tests (Figure 14). co-IZmo/D integrates air conditioning features and IT equipment modules into an ISO-standard 20-foot container without the use of external air conditioners or humidifiers. In one of our tests, dubbed the 2013 IT Equipment Adaptation Test, we worked with major IT equipment vendors to evaluate the characteristics of IT equipment in summer and winter in a completely outside-air air conditioned environment and assessed the health of IT equipment under year-round outside-air cooling.

We ran the tests on IT equipment supplied with a maximum air temperature of 35°C/40°C, and the IT equipment functioned normally in a high-temperature environment of 45°C without any drop in CPU processing performance etc. In a high-humidity environment, however, we did observe a noticeable deterioration in the IT equipment boards. From these tests, we concluded that high humidity would reduce IT equipment operating life and increase fault rates. The tests also revealed that the rotation speed characteristics of IT equipment internal fans vary a bit depending on the

Figure 14: co-IZmo/D

IT equipment manufacturer's design (Figure 15). From a facilities perspective, letting server room temperature rise by 1°C will save on air conditioner power consumption. But with tens of thousands of IT equipment units all running in a unique way as the temperature rises, the power consumption of those units also changes. To maintain data center quality and achieve true energy savings, therefore, we determined that, instead of simply raising server room temperature, we should understand the individual characteristics of the IT equipment and control operations to be jointly optimal for both the IT equipment and the facility itself.

AI Control

Shirai DCC uses AI-based air conditioning control that minimizes the combined power consumption of the IT equipment and air conditioning systems, the main data center elements, based on what we learned from our testing efforts, including chillerless operation tests and IT equipment adaptation tests. IT equipment characteristics such as rated current and current draw at different temperatures (as measured in Shiroi DCC's test environment) are recorded in a database, and measurements of server room temperature, actual IT equipment current draw, amount of heat processed, and so forth are taken and analyzed while the facility is in operation to provide (rule-based) estimates of the characteristics of each individual rack. Similarly, air conditioning system characteristics are estimated based on certain operating conditions, parameters, and so forth, and a rule engine infers the operational settings that will minimize the combined power consumption of the IT equipment and air conditioning systems, with the result being to optimize air conditioning system operating conditions.

So we have deployed the direct outside-air cooling systems developed at Matsue on system modules that minimize pressure loss by virtue of a large space, large apertures, and wide span, and combined this with AI control, and we intend to proactively adopt new technologies going forward, our ultimate goal being to achieve an annual average PUE of under 1.2.

Lithium-ion Batteries

We have chased energy savings at Matsue DCP through the use of outside-air cooling and air conditioning units with high COP scores, and by selecting high-efficiency IT

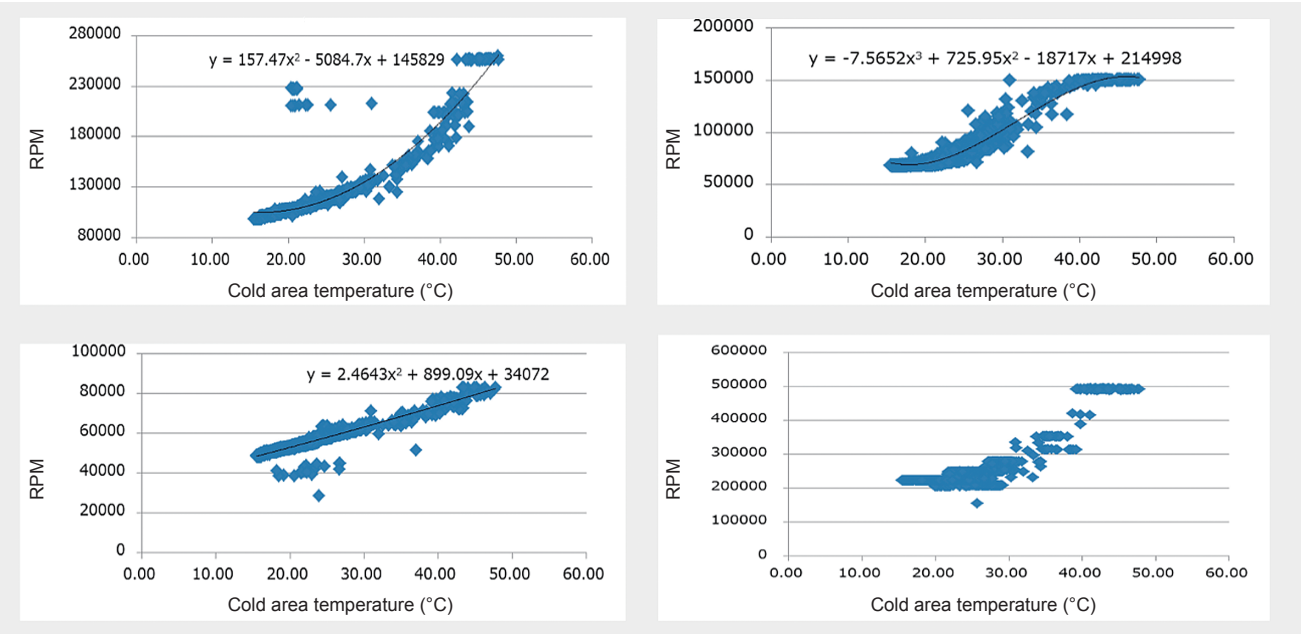

Figure 15: IT Equipment Internal Fan Rotation Speed Characteristics by Manufacturer (Chillerless Tests, 2012)

equipment. Air conditioning systems and IT equipment are always the main factor in data center energy savings. The key element in a data center, however, is the power system. We see the power system as the next key avenue for saving energy. Accordingly, we adopted a three-phase four-wire UPS and bus ducts at Matsue DCP Site 2, and we also continue to evaluate software for predicting electricity usage and for peak-shaving control, and continue to devise and evaluate mechanisms for the selective supply of power from fuel cells, PV systems, and DC UPS systems.

Before designing Shiroi DCC, we also conducted a technology survey, mainly looking at lithium-ion batteries, with a focus on two particular themes: (i) highly economically viable power generation systems based on summer usage, peak usage, etc. and (ii) electricity/power systems capable of delivering high efficiency in view of future technology trends. Based on this, we selected Powerpack, a lithium-ion battery storage solution from Tesla that offers operating control features, installing it in November 2019 (Figure 16). Powerpack replaced the lead-acid battery UPS that we initially intended to install as a backup power supply for the air conditioning systems, and it offers the added feature of peak shaving at a similar cost level to the lead-acid battery UPS. Powerpack enables peak shaving and load shifting with respect to the power received by the data center, which peaks during summer daytimes. It makes it possible to reduce basic charges by reducing the demanded power and to reduce electricity charges by using power at off-peak times and thus purchasing electricity at lower prices (Figure 17).

In the past, UPS systems, emergency generators, and other such power systems played a behind-the-scenes role in ensuring high quality. At Shiroi DCC, we are working to develop a new business model by actively looking for more ways to use such systems as data center energy resources.

3.4 Carbon Neutral Data Center Model

■ IIJ's Carbon Status and Road to Carbon Neutrality

Manufacturing industries are being called on to reduce the CO2 emissions that result from making products in factories, and the transport industry to reduce the CO2 emissions from trucks. Given that over 90% of IIJ's CO2 can be traced back to the power consumed by its data centers, the use of energy-saving equipment and renewables

Figure 16: Powerpack installed at Shiroi DCC

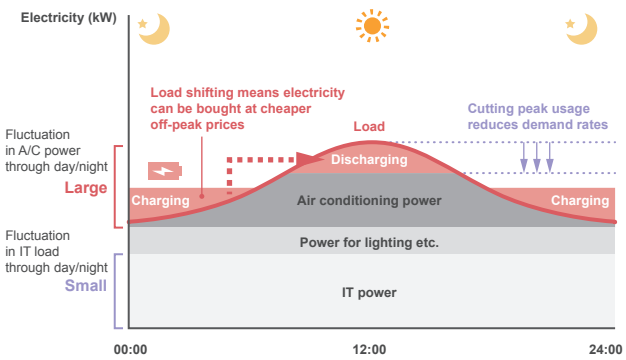

Figure 17: Data Center Power Demand and Peak Shaving / Peak Load Shifting

directly impacts on the carbon neutrality status of IIJ as a whole.

We plan to continuously improve the energy-saving technologies implemented at Matsue DCP and Shiroi DCC as described above to make the facilities even more energy efficient. Meanwhile, we have only just started working on the use of renewable energy, and one of the early possibilities for tackling this will be to switch to an electricity retailer supply contract for the supply of renewable energy certified as non-fossil fuel electricity. There are disadvantages, however. Renewable energy is more expensive than ordinary electricity, and there is no guarantee that it will be supplied at a stable price over the long term. So with the cost of generating renewable energy falling year by year, possible next steps may be to buy electricity directly from renewable energy plants or to have our own generation plants.

Although renewable generation costs are falling, power generated from non-solar sources including wind and biomass remains more expensive than ordinary power purchased from electric utilities. And building a power plant will take time, including the time needed to obtain a suitable site, so it is more realistic as a medium- to long-term proposition. Near term, we need to implement mechanisms for the supply of power with a focus on solar, or photovoltaic (PV), generation.

Carbon Neutral Data Center Model

In light of all this, with a view to achieving carbon neutrality, we have defined a carbon-neutral data center reference model with the features listed in Table 1, and we plan to continue modifying our data centers and building new ones in this light.

No.	Feature
1	Use of energy-saving technologies, such as outside-air cooling systems that substantially reduce air conditioning power consumption and an efficient three-phase four-wire UPS, to reduce absolute power consumption levels.
2	Absence of other equipment and fixtures on building roofs to make large roof areas available for the installation of on-site PV systems.
3	In the short term, procure power from a combination of extra-high-voltage, high-voltage, and low-voltage off-site PV systems. Long term, allow for the possibility of procuring power from other sources such as wind and hydrogen.
4	When daytime power supply exceeds consumption, store the excess in batteries for use at night.
5	Use of network measurement and control functions between generation equipment and the data center to balance the amount of power generated and the amount consumed.
6	Use of IT load control to channel loads toward times of surplus power.

Table 1: Features of a Carbon-Neutral Data Center

Type of facility	Office building	Warehouse	Conventional data center	Carbon-neutral data center
Power consumption	Low	Low	Medium to high	Medium to high
Power density	Low	Low	High	High
On-site PV installation space	Small Typically, external A/C units etc. are installed on rooftops	Medium Large-scale rooftop installations possible	Small to medium Similar to office buildings, but dedicated equipment can be used to create space	Medium Roof structure allows for large-scale rooftop installations
Typical daily power profile Generated Consumed	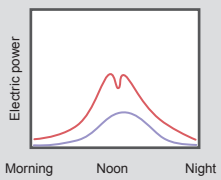	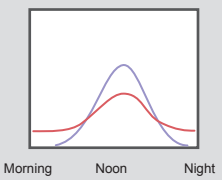	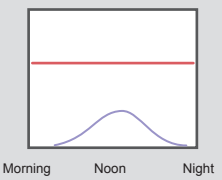	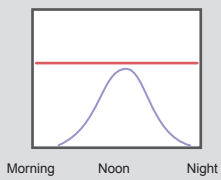

Table 2: Differences in Electric Power by Type of Facility

The first feature is the use of energy-saving technology to the greatest possible extent to reduce absolute power consumption.

The second feature is the installation of as much on-site PV generation capacity as possible, as the cost of generating electricity this way is now lower than ordinary electricity charges. As Table 2 shows, data centers bring together vast amounts of IT equipment and thus consume several dozen times more power per unit floor area than office buildings and the like (said to be around 50–100W/m^2). As they consume so much in absolute terms, PV generation alone cannot keep up with demand for power across the entire data center. But by adopting a warehouse-like building structure and making the roof, where PV systems are installed, as large as possible, you can increase low-cost on-site PV generation capacity.

The third feature is the installation of off-site power generation equipment (off-site equipment that supplies power to your site via the power utility grid). As Table 3 shows, off-site PV is more expensive than on-site PV, but increasing off-site PV capacity is a viable option for securing the required amount of power.

Off-site PV systems can be classified into three types according to the voltage at which they connect to the power grid, as shown in Table 4. To adequately supply a several-dozen-MW-scale data center, a large, extra-high-voltage PV system is efficient, but these tend to take a long time to build because you have to find a suitably large site, negotiate connection terms and conditions with the electric utility, and so forth. The high-voltage and low-voltage options are quicker to build, but if you want to generate the same amount of power as a 50MW extra-high-voltage PV installation, for example, you will need 1,000 50KW low-voltage generation plants. Given the amount of power required and time required, we think off-site PV power needs to be sourced from a combination of extra-high-voltage, high-voltage, and low-voltage systems.

System type	On-site PV	Off-site PV
Characteristics	・No need for a dedicated site if installed on your building rooftops etc. ・In contrast with off-site PV, involves no wheeling charges or renewable energy levies, and use of subsidies can make it even more cost effective. ・But scale is small (can only supply a few percent of total data center usage).	・Initial costs, including those for the land, are higher than for on-site if you take on the capital investment yourself. ・Need to enter into long-term contracts if having power supplied from another company's equipment. ・Involves wheeling costs as power is supplied via the power utility grid. ・Systems/procedures can change. ・Can adjust the amount of power generated by increasing scale or adding plants.

Table 3: Differences Between On-site and Off-site PV

Power grid connection voltage	Extra-high-voltage 2,000kW or more	High-voltage 500–2,000kW	Low-voltage up to 50kW
Required site area (at 10m^2/KW)	20,000m^2 (2ha) or more	500m2 – 20,000m^2 (2ha)	up to 500m^2
Initial investment	22.2 yen/kW 440mn yen or more	22.2–25.5 yen/kW 13–440mn yen	25.5 yen/kW up to 13mn yen
Construction period	1 year or more	up to 1 year	up to 6 months
No. plants needed for 50MW output	5MW/plant 10 plants	500kW/plant 100 plants	50kW/plant 1,000 plants

Table 4: PV System Categories

Storage is the fourth feature required for carbon-neutral data centers. As Figure 18 illustrates, as off-site PV increases, the power supplied by PV capacity in the daytime can overshoot data center power consumption, in which case it is stored and used at night. This makes it possible to use renewable energy at night as well, but the fact that the cost of storage batteries is still high is an obstacle to the full-fledged rollout of such systems.

The fifth feature is the use of network functions that enable measurement and control of what happens between the power generation equipment and the data center. This makes it possible to balance supply and demand and ensure the stable operation of the grid.

The sixth feature is the use of IT load control, which involves, for example, halting servers on nights when there is not enough power and only running them in the day, so that the load is channeled toward times when there is surplus power. Information systems generally cannot be stopped for 24 hours at a time, but we can expect more and more new use cases amenable to load control to appear—data analysis tasks that can be given a manageable timeframe to complete or mining tasks that require low-cost processing.

Figure 19 illustrates how a carbon-neutral data center with these features could be set up. To realize carbon neutrality while maintaining the reliability and other quality attributes required of a data center, we need to create a new model

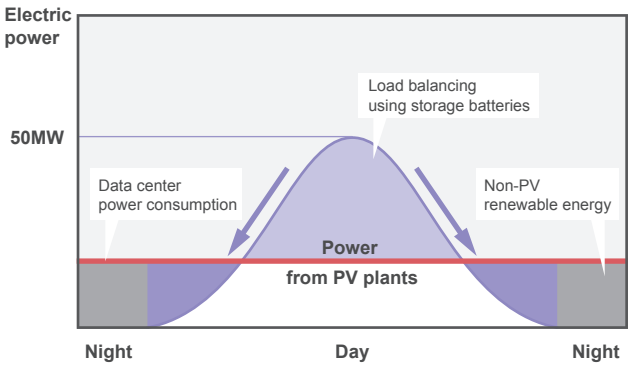

Figure 18: Data Center Power Supply/Demand Overview

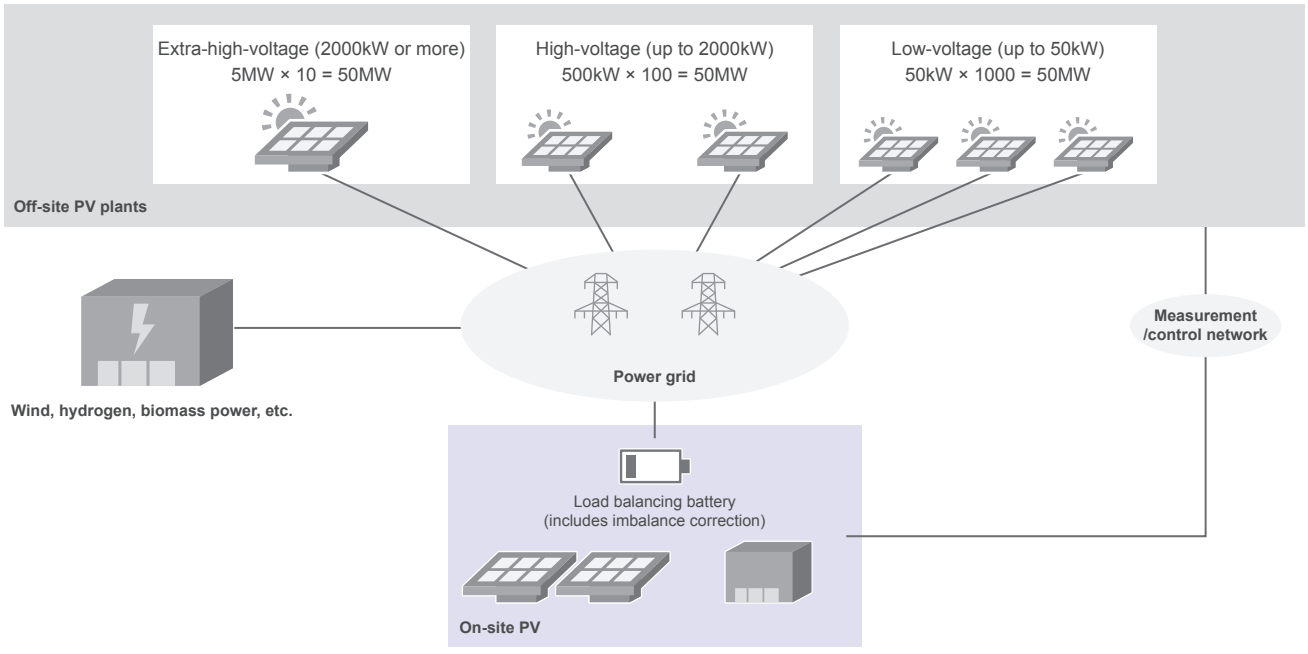

Figure 19: Carbon-Neutral Data Center Model

that organically combines the generation equipment that supplies the power and the data center that consumes that power. While we only have a conceptual model to illuminate the way forward at this stage, we will be conducting technical tests and working with external partners going forward, on both the business and technological fronts, to sort out the details and apply what we have learned when modifying our own data centers or building new ones.

3.5 Conclusion

IIJ has also been involved in external initiatives using the technologies it has developed in-house, one example being a NEDO (New Energy and Industrial Technology Development Organization) demonstration project aimed at using the JCM (Joint Crediting Mechanism) to reduce greenhouse gas emissions, as part of which IIJ provided support for the construction and operation of a highly energy-efficient containerized data center in the Laotian capital of Vientiane.

The road to carbon neutrality will no doubt be long and challenging, but we will continue to take advantage of Internet technologies, including IoT and mobile, and test and implement systems, and we will seek to apply the knowledge and insight we gain not just internally at IIJ but also in the world at large, our aim being to contribute toward carbon neutrality in the data center and, ultimately, across society as a whole.

Isao Kubo

General Manager, Infrastructure Services Department, Infrastructure Engineering Division, IIJ.
Mr. Kubo joined IIJ in 2008. He oversees the data center business and the construction of Matsue DCP and Shiroi DCC. His aim is to achieve carbon neutrality as soon as possible.

3.2 Outcomes at Matsue Data Center Park
Yoshimasa Kano

Manager, Matsue Data Center Park, Infrastructure Services Department, Infrastructure Engineering Division, IIJ.
Mr. Kano joined IIJ in 2016. He is engaged in the operation of containerized data centers and the testing of next-generation modular data centers.

3.3 Initiatives at Shiroi Data Center Campus
Akio Hashimoto

Deputy General Manager, Infrastructure Services Department, Infrastructure Engineering Division, IIJ.
Director, Shiroi Data Center Campus, Infrastructure Services Department, Infrastructure Engineering Division, IIJ.
Mr. Hashimoto joined IIJ in 2009. He works on the study and design, construction, and operation of next-generation data centers and actively pursues data center automation and efficiency gains.

IIJ's Road to BCR Approval —Complying with EU GDPR and Beyond

In August 2021, we received notice of the long-awaited approval for IIJ's BCRs (Binding Corporate Rules) from LDI-NRW, the supervisory authority in North Rhine-Westphalia (NRW), Germany. Since the EU's GDPR (General Data Protection Regulation) came into force, 18 companies worldwide have received BCR approval (as of August 2021), IIJ being one of them. The five years since we started looking at BCRs as a means for complying with the GDPR have been eventful, and with the approval, we can now proudly say that our efforts so far have finally been recognized.

Here, I discuss why we chose the IIJ BCR approval route, what the road to acquiring that approval was like, and where we are headed from here.

4.1 Decision to Seek BCR Approval

The EU drafted the GDPR for the purpose of protecting personal data within the EEA (European Economic Area) in 2016, and it came into effect in 2018. The GDPR was the subject of a lot of speculation back when it was announced, and I remember some extreme arguments saying, for example. that it would harm the free distribution of information on the Internet.

The EU's objective was of course not to restrict the distribution of personal data on the Internet. Rather, it was calling for the appropriate use of data on people in the EU at a time when such information was already being used in countries everywhere. Where practices had been vague, the EU wanted to lay out clearer protections on personal data that had real effect. It was a truly advanced initiative to institutionalize such protections. I think that, even now, the policies and practices for personal data protection laid out in the GDPR have a lot to say about personal data protection in the information age.

With the rise of major platform operators, exemplified by GAFA, that use personal data to the hilt and Snowden's exposure of the US government's excessive surveillance practices, there was strong distrust within Europe about the US's use of personal data, and I think this is the real backdrop for what prompted the creation of the GDPR. I think, at its heart, this traces back to a marked difference in cultural values about personal data. This is something I will discuss later.

It became evident that the EU was pretty serious, but IIJ doesn't provide personal services on a global basis, and there was talk about the Japanese government's Personal Information Protection Commission working toward some sort of national certification from the EU anyway, so truth be told, we hadn't intended to actively pursue GDPR compliance. But as we worked through all the factors to consider, and with a proposal from Shinpei Ogawa, a director of IIJ Europe at the time, we came to the conclusion that seeking IIJ's own BCR approval would be advantageous.

We had four main reasons for embarking on our journey for IIJ's BCR approval.

■ We saw the need for global security governance

IIJ has been deploying group companies all over the world for some time, and doing business naturally requires us to address personal data protection and other factors in each country. In the past, however, we did not fully comprehend the situation in every country, so it was left up to each group company to mount its own response, and to be honest, we weren't able to address the issue of control on a group level. Creating IIJ's BCRs to comply with the GDPR and putting this into effect at group companies was crucial in enabling

us to implement control at a time when the situation around personal data was becoming more and more nuanced.

Alternatives to BCR were too complicated

An alternative to obtaining BCR approval as a means of meeting the GDPR requirements is to enter into SCC (Standard Contractual Clauses), a form of template contract, as necessary and perform data transfers on that basis. These contracts, however, must be meshed in the sense that there needs to be contracts between the controller (the party responsible for holding the personal data) and the processors (the parties who receive the personal data for processing), as well as between processors themselves. And if changes are made to how the data are processed, the contracts must be entered into again, so managing the contracts is cumbersome. It was easy to see that we would end up with quite a lot of contracts just with IIJ's own group companies, and given the nature of our business, in many cases information from our customers is deployed across each of our group companies' platforms. For these reasons, we were concerned about this approach eventually collapsing. IIJ's own BCR approval was really the only way to solve the problem.

GDPR represented a more advanced approach than Japan's own personal information protection practices

IIJ naturally implements internal controls for the protection of personal information and has acquired Japan's P Mark (Privacy Mark). Yet something that became eminently apparent as we started to understand more about the GDPR is that the EU's commitment is on another level. Japan's personal information protections certainly do comprise a range of carefully thought out measures, but the EU's approach was to tackle these difficult problems through sound reasoning and ideas, and to look ahead and start implementing responses to a range of issues that potentially lay down the track. We were honestly surprised at how far they were taking it. Japan was also naturally making an effort in that regard, and the EU did adopt an adequacy decision on Japan, but it is telling that supplementary rules were applied to transfers of personal data from within the EEA to Japan. This leads into the discussion that follows, but the approach of giving careful thought not just to the controller but also to the processors, and developing a framework with respect to the processors, is a fairly advanced one. The EU obviously implemented this because it felt it was necessary to protect personal data, but I think it also shows that they were somewhat prescient about how personal data would be handled in the Internet's cloud era.

It would enable protection of personal data on cloud services in addition to personal data held by IIJ

The EU's motivation for creating the GDPR was clearly to regulate the handling of personal data on the Internet by platform operators and the like within the EEA, to which end it sought to impose rules on business operators and the like that handle personal data so as to ensure that they do so properly in accord with EU standards. As such, the GDPR naturally applies not only to controllers responsible for managing personal data but also to processors who receive such data from controllers. I think the EU made such provisions explicit because it was seeking to confront what it saw as a less than appropriate state of affairs particularly with respect to the platform operators who process personal data. This potentially opened the way for outsourcing, such as cloud services, to comply with the GDPR by meeting the BCR conditions under the processor category. That is, it provided a way for controllers, the customers of such services, to prove that they use processors who have implemented proper protections.

IIJ Europe began looking seriously at GDPR compliance in January 2016, and IIJ's board of directors officially gave approval to start working toward BCR approval in May 2016. The approval was actually received on August 5, 2021, partly because of how long the review takes, but also due to the impact of Brexit and other subsequent events. While not all that many people were involved over that period, a number of internal units, led by the Risk Management Office, had a role to play, including the Business Risk Consulting Headquarters, the Global Business Division, and the Compliance Department, and we actually used a law firm when negotiating with the EU's supervisory bodies as well. Our efforts did finally pay off in the form of BCR approval. IIJ is the first global cloud vendor, not just in Japan but worldwide, to receive approval since the GDPR took effect.

But as with other aspects of data governance, obtaining approval is not the goal. Instead, it provides an opportunity to embark on new efforts to establish internal control for global data governance.

4.2 What are BCRs?

Below is a brief explanation of BCRs.

BCRs are data protection policies that are adhered to by an entire corporate group (these policies are also made known to the subjects of personal data and are thus widely disclosed) and embody rules that are binding on the group companies and their employees. The objective is data protection, but this is not so much about implementing technical security measures as it is about sharing information and educating employees on the basic principles and rules around the handling of personal data, creating an operating environment that allows employees to raise queries or complaints, and so forth. In that sense, the approach is akin to the efforts developed in Japan over the past many years around internal control for information security and personal information protection. That said, the EU's GDPR is the most stringent personal information protection regulation anywhere in the world, so obtaining BCR approval is by no means easy.

At the IIJ Group level, the IIJ Binding Corporate Rules appear at the very bottom of the privacy policy on IIJ Europe's website, and IIJ Group companies also link to this.

Once a corporate group's BCRs have been approved by the competent data protection authority in the EU, this certifies that, in terms of personal data protection under EU law, the group has appropriate safeguards in place, and under GDPR Article 46, Item 2(b), this allows it to legally transfer personal data fro within the EEA to locations outside of the EU.

Transfers of personal data from within the EEA to the outside of the EEA are in principle prohibited. There are a number of recognized ways of making it possible to do this, though. Among them, BCRs represent the strictest standards at the corporate level. As long as they comply with the BCRs, transfers of personal data outside the region to corporate groups that have BCR approval, or between companies within a group that has received approval, are recognized as being subject to the appropriate safeguards required by the GDPR.

Approval is obtained by submitting BCRs as stipulated in GDPR Article 47 to the competent data protection authority for approval. A rigorous review process that goes beyond the competent authority ensues. First, the competent authority reviews the BCRs with the assistance of supervisory authorities, and corrections are repeatedly made

in conjunction with the company under review. Once it is satisfied, the competent authority communicates with the EDPB (European Data Protection Board), part of the European Commission, the EU's executive branch, and submits a pre-review request to the ITES (International Transfer Expert Subgroup) meeting. ITES is made up of experts from the supervisory authorities of all EU member states. They may think that something should be done differently or that a particular rule is a little loose, and they accordingly send revision requests to the competent authority, based on which the corporate group revises the BCR draft. This process is repeated until it is apparent there are no further opinions on the pre-review, at which point the competent authority asks for an Opinion from a plenary meeting of the EDPB, the highest decision-making authority in this case, composed of representatives of all EU national data protection authorities. The Opinion is an official document of the EDPB, in accord with which the competent authority provides instructions to the corporate group to finalize the BCRs, and then grants final approval. This rigorous process means that obtaining BCR approval from your competent data protection authority in the EU takes quite a lot of effort and time.

In the IIJ Group's case, we submitted our BCRs to the UK's ICO (Information Commissioner's Office), our competent authority at the time, in October 2016. The impact of Brexit, however, means that the competent authority for us is now the supervisory authority in North Rhine-Westphalia (NRW), Germany. IIJ Europe (based in London) had been the IIJ Group's headquarters in the EU, but the UK's withdrawal from the EU meant that this role passed to IIJ Deutschland (based in Dusseldorf), and thus the supervisory authority in NRW, in which Dusseldorf is located, became the IIJ Group's competent authority. The norm in other countries is to have a single national authority, but with Germany being a federation, all 16 of its states have their own personal data protection supervisory authority.

The EU adopted an adequacy decision on Japan on January 23, 2019. The decision means the EU recognizes that Japan has personal information protection guarantees that are in line with those that apply in the EEA. This recognition from the European Commission was the result of work by Japan's Personal Information Protection Commission along with other stakeholders and no doubt came as a boon for many Japanese companies.

It does not mean, however, that all Japanese companies are now GDPR compliant.

First of all, the GDPR does not in principle allow the transfer of personal data out of the EEA, so the cross-border transfer rules must be observed when transferring data. The cross-border transfer rules, per GDPR Articles 45 and 46, allow such transfers under conditions including the following.

- The European Commission has adopted an adequacy decision on the country
- BCR approval has been acquired
- SCCs have been entered into
- In compliance with a (an industry) code of conduct approved by the EDPB

Put differently, only transfers of personal data out of the EEA pursuant to the GDPR are exempted. In addition, the Personal Information Protection Commission has published official notice that supplementary rules will apply to personal data transferred from within the EEA wherever the local rules are deemed inadequate. The EU's adequacy decision means that transfers of personal data from the EU

to Japan are permitted, but such data may not be further transferred from Japan to a third country. So the adequacy decision does not fully cover cases in which, for instance, an employee register is shared globally.

Japan's pursuit of an adequacy decision was the right approach for the nation to take, and an understandable one, in view of the circumstances of many Japanese companies' businesses, but it is unfortunate that it led to a sense that companies now no longer need to take any particular steps of their own with respect to the GDPR. Personal data protection is an extremely important issue for the Internet and other types of new information infrastructure, and we should not forget that the situation these days is such that even companies that do business mainly within Japan cannot ignore its impact.

4.3 Personal Data Protection Initiatives Around the World

As my mind became increasingly wound up in the various personal data protection initiatives out there, I was made acutely aware of differences in perspectives on personal data and astonished by just how different the cultural backgrounds can be. Norms around personal data in Japan merely represented Japan's own local perspective. My realization that there are rights and responsibilities with respect to personal data in other parts of the world that differ completely from those found in Japan was a very important insight that I gained from IIJ's BCR approval process. To be honest, I realized that IIJ's mindset as a traditional Japanese company simply would not cut it on the world stage. There exist multiple ideas about what is right and just when it comes to personal data around the world, and they will all no doubt have a major influence on the information society we live in going forward.

I will now briefly explain the background to the creation of the GDPR. An essential element is that the EU authorities are in opposition to the US in terms of privacy protection.

■ Historical background

The 9/11 terrorist attacks in the US were more than enough to shake the world's collective consciousness. After 9/11, the US embarked on a secret mass surveillance program and implemented mechanisms for installing backdoors in social media and other services in order to expose terrorists. These operations should only have been carried out under a court order, but in reality, as exposed by Edward Snowden in June 2013, the US had made it possible for intelligence personnel to freely snoop into people's privacy in the course of their intelligence activities. I think it was difficult for the American people to oppose this because the atmosphere was such that they felt compelled to allow the government to do what it needed to do for national security. I think Americans have a very strong sense of ownership about their country in the sense that they see the government as the people's representative and spokesperson, an attitude that dates back to the nation's founding. But then Snowden revealed that the US had been eavesdropping indiscriminately on embassies of the country's allies including Germany and France. This enraged the EU's member states and had major ramifications, the German government's axing of its Verizon contract, for instance.

The idea that the government is absolutely right does not exist in the EU. For more than a millennium, the European continent was again and again the scene of wars and conflicts that ultimately failed to settle national borders or territory. The rise of tyrants repeatedly led to tragic events—people oppressing those from different denominations within their own religion, for example, or

massacring other ethnic groups. The European Coal and Steel Community was created with the aim of establishing final national boundaries and securing a lasting peace after World War II. Gradually building its cooperation with the European Atomic Energy Community and the European Commission, it eventually grew to become the EU we know today. The EU operates on the premise that governments can also overstep and run wild, so when EU laws are passed, it is stipulated that independent, third-party supervisory authorities that have the power to enforce the law with respect to governments also be established. So the EU's privacy protection supervisory authorities strictly enforce the law with respect to privacy breaches not only over private-sector companies but also over governments and public agencies.

This difference in thinking means the US and the EU are basically incompatible in the world of privacy. But because they are each other's largest trading partner, people in the EU involved in the business of trade desire a good relationship with the US. Personal data is generally distributed as part of commercial activities and is thus inseparably linked with trade negotiations. In 2000, therefore, the EU created a framework called the EU-US Safe Harbor Principles, under which private-sector organizations were permitted to freely transfer personal data from the EU to the US provided they submitted a notification to the US Department of Commerce attesting that they had adequate security measures in place and will not pass personal data from the EU to third parties.

This collapsed in the wake of the Snowden expose. Snowden's revelations prompted doubt about the effectiveness of the Safe Harbor agreements in the mind of Austrian lawyer Max Schrems, who then filed a legal complaint in Ireland. The details are complex, but roughly speaking, under Facebook's Safe Harbor agreement, Facebook users' personal data was, for example, being transferred from Ireland (location of Facebook's EU headquarters) to the US. But as the Snowden expose revealed, the US government was able to freely view that data. This means, said the complaint, that Facebook is breaking the Safe Harbor agreement, and if the US government can legally spy on the services of private companies under US law, then the Safe Harbor agreement itself is pretty much meaningless in the first place. The Irish High Court, unable to make a decision, referred the matter to the Court of Justice of the EU (CJEU). In a shocking ruling, in October 2015 the CJEU declared the Safe Harbor Agreement invalid. This is known as the Schrems I decision. The inability to transfer personal data from the EU to the US represented an extreme impediment to trade, however, so economic proponents within the European Commission and the US Department of Commerce adopted a new special framework called Privacy Shield in August 2016. The Patriot Act, created after the 2001 terrorist attacks, expired in 2015, but the mechanisms for exposing terrorists arguably remained intact, albeit with increased transparency, under its successor, the Freedom Act. Privacy protection advocates within the European Commission claimed, therefore, that nothing had really changed, and debate about the validity of Privacy Shield thus continued to smolder on during and after 2017. Once the GDPR came into effect on May 25, 2018, Schrems immediately filed another suit claiming that it was illegal for US companies to transfer EU personal data to the US based on Privacy Shield. The case was settled in July 2020 with the ruling that Privacy Shield was also illegal. This is the Schrems II decision. Following Schrems II, the SCCs were also reviewed, and modernized SCCs were issued in June 2021.

The new SCCs provide extra protections that require companies and other bodies that transfer data to disclose information about access by public authorities in the destination country (whether the country has laws that allow the government requisition data and whether they are actually enforced).

Once personal data is transferred from one country to another, it is no longer protected by the origin country's laws, so various restrictions are thus imposed. Meanwhile, the EU does recognize a number of countries as offering adequately secure personal data protections. It refers to these as "adequate" countries (recognized by adequacy decisions), and they include Switzerland, New Zealand, and Argentina. Personal data can be transferred from the EU to these countries. Japan joined these ranks in January 2019. Japan's Personal Information Protection Act is not quite up to the GDPR standards, however, so personal data may only be transferred if certain supplementary rules are applied.

The Personal Information Protection Act is to be reviewed every three years to bring it up to speed with technological developments and other countries' laws and regulations, and thus it may approach parity with the GDPR going forward. On the other hand, Japanese personal data can also be transferred to the EU. As announced in January 2019, Japan and the EU reached a mutual agreement on the flow of personal data between each other's domains, and this means that Japan also recognizes the EU as having adequate protections. The only other place Japan recognizes in this manner is the UK. Given these developments, it does seem like the Japanese government's approach when it comes to protecting privacy is to strengthen individuals' rights and interests using the EU's GDPR as a reference.

Although the process is gradual, the revised Personal Information Protection Act, which was passed in June 2020 and will come into effect from April 1, 2022, also looks set to tighten Japan's protections and bring them closer in line with the GDPR. For instance, penalties will be raised from 300,000 yen and 500,000 yen to 100 million yen. Stronger information disclosure requirements will apply when personal data is provided to a third party in a foreign country (from the EU's perspective, this means when personal data is transferred out of its domain). Cookies will be subject to restrictions, although not quite to the extent as in the EU. And it will be mandatory to report personal data leaks that match certain patterns to the Personal Information Protection Commission and other bodies (preliminary report within 3–5 days and a final report within 30 or 60 days; incidentally, the GDPR mandates reporting within 72 hours).

Cultural background

In light of the above historical background, I see three major trends in personal data protection across the globe today.

One is a sort of public welfare idea that the sharing of big data will benefit people all over the world, which seems to be the mainstream view in the US. Another is the human rights assertion that says that the ability to manage your own personal data is a basic human right, which is the mainstream view in Europe. And then there is the security-oriented idea that the management of data within a country's own domain is its own national security issue, which is probably exemplified by China.

In the US, personal data is seen not only as holding convenience for the user but also as conveying benefits to a wide swath of other users, as exemplified by the rise of GAFA. This is evident from Google's mission, for instance,

which reads: "Our mission is to organise the world's information and make it universally accessible and useful." The history of American society is one of pioneering on the frontier, and this is possibly what embedded the idea of information sharing, including for the purposes of protecting national security, as an important part of its culture. In terms of its development strategy too, this seems to have led to a set of values that embrace the idea of de facto standards, the idea that prevalence itself is what makes a standard, something that has been a prominent factor in the Internet as well.

In Europe, on the other hand, personal data was at times something that could affect whether a person lived or died, so the ability of the subject of personal data to know and manage that data is seen as a right. Further, the development strategy is one of clarifying processes thought to have been developed as part of colonial policies, with standardization being pivotal to the global development of data protections.

In China, which exerts a significant influence on the world in recent times, the desire to maintain the nation is a crucial factor, and so naturally the thinking is that judgements about all sorts of information, including personal data, should be based on national security. Hence, China recommends its own country's services over those created in the US, and it imposes heavy restrictions on information flows between countries. And of course, the strategy is a nationalistic one: information flows involving China will take place under China's restrictions.

Of course, these are somewhat stereotypical characterizations. I think countries' information control strategies are carried out with an eye on any number of elements to do with the public interest, rights, and security, rather than emphasizing any single factor alone. But I think it's important to note that, broadly speaking, there are three major perspectives around personal data, each emerging from its own historical and cultural backdrop, and there is no unified mindset on personal data that applies across the entire world.

Closer to home in Japan (and while certainly not the mainstream world view), there is a culture of regarding information as being equivalent to value itself, a view that comes from the Japanese concept of kotodama, the belief that mystical powers dwell in words. So when it comes to personal data, people may feel that if someone knows information about you, they actually know you in some true sense. In older times, there was a cultural practice of hiding one's true name out of the belief that knowing someone's name would confer dominion over them. There is also a long-held idea that you should not utter unlucky thoughts (put misfortunes into words) because expressing information can somehow lead to the events it describes coming to pass (this is a deep rabbit hole to go down, so I will cut the discussion short here).

In any case, when it comes to personal data protection, this Japanese style of thinking, at least, is unlikely to pass muster elsewhere in terms of the mindset and meaning it implies and the weight of the belief. Of course, there is clearly no globally unified view, and this holds for the US mindset, the European mindset, and the Chinese mindset alike. In our information society, information is not just a set of symbols. There are historical and cultural backgrounds, and therefore each culture has its own principles, its own ideas about the right way to handle information, so when it comes to flows of information across national borders, there is a desire to ensure that the principles of each country or region involved are respected.

4.4 IIJ's Road to BCR Approval

IIJ was an early mover on BCR approval, but the process was full of twists and turns. The impact of Brexit in particular was significant, and the IIJ Group, which had aimed to obtain approval from the UKs ICO, did have to make a major strategic change along the way.

Table 1 provides a brief overview.

4.5 Looking Ahead

As discussed earlier, as the Internet spreads further, initiatives in personal data protection will no doubt continue to affect Internet-related companies and other companies in various ways. This process is a necessary part of the way society is adapting to facets of the information age, most notably the Internet, and even in Japan where attitudes can be indifferent, these changes are something that, sooner or later, we will be unable to ignore. Of course, even from my own experience, I cannot declare that every company should obtain BCR approval, but objectively speaking, there are many companies in Japan that do need it, and they will eventually be forced to take action of some sort.

Date	Organization	Action
Jan 2016	IIJ Europe	IIJ-EU receives CEO approval to start working toward GDPR compliance. Work begins in earnest.
Mar 2016	GDPR Office	Office set up by the Risk Management Office, Compliance Department, Global Business Division, and IIJ Europe.
Jun 2016	UK	Brexit prevails.
Jul 2016	GDPR Office	Decides on law firm.
Aug–Oct 2016	Risk Management Office	Creates BCR document.
Aug 2016	IIJ Europe	Starts GDPR compliance support consulting.
Oct 2016	GDPR Office	Submits BCRs to the Information Commissioner's Office (ICO), the UK's personal data protection authority.
Mar 2017	UK	Prime Minister signs a letter triggering Brexit.
Aug 2017	ICO	First communication that IIJ's BCRs are now under review. An ongoing process of revisions follows.
May 25, 2018	EU	GDPR comes into effect.
Jan 10, 2019	ICO	UK ICO's review is completed and submitted to the co-reviewers (Germany, Netherlands).
Feb 12, 2019	EDPB	Information on competent supervisory authorities post-Brexit is published. Appears we will not make it in time for Brexit.
Mar 1, 2019	LDI-NRW	Co-reviewer comments on IIJ's BCRs from the authority in NRW, Germany.
Mar 21, 2019	Dutch authority	Co-reviewer comments on IIJ's BCRs from the Dutch authority.
Mar 28, 2019	GDPR Office	Presents response to comments to ICO. Brexit deadline extended to Oct 31.
May 16, 2019	ICO	Co-reviewer provides notice that the review is complete. Pre-review by the EDPB ITES (International Transfers Experts Subgroup) meeting requested.
Jun 2019	UK	Prime Minister Theresa May resigns.
Jul 2019	UK	Boris Johnson becomes Prime Minister.
Jan 31, 2020	EU/UK	Britain withdraws from the EU. The transition period runs until end-2020.
Apr 2020	EDPB	Pre-review conducted by the EDPB ITES meeting under the ICO's guidance. Several rounds of revisions ensue.
Jun 2020	EDPB	Provides notice that ICO approvals made during the Brexit transition period are invalid.
Jul–Sep 2020	GDPR Office	The supervisory authority in NRW, Germany, where IIJ Deutschland is located, is appointed as the competent authority.
Sep 2020	GDPR Office	Formally requests the NRW authority to act as the competent authority.
Sep–Nov 2020	ICO/LDI-NRW	LDI-NRW takes over from ICO on IIJ's BCRs.
Dec 2, 2020	LDI-NRW	LDI-NRW officially becomes IIJ's competent authority.
Dec 31, 2020	UK	Completely withdraws from the EU.
Apr 2021	EDPB	Review at ITES meeting (review had progressed under ICO, so passage is smooth).
May 2021	GDPR Office	Creates a German-language version of the BCR document.
Jun 28, 2021	LDI-NRW	Submits BCR-C and BCR-P to the EDPB.
Jul 28, 2021	EDPB	Deliberates at a plenary meeting. All clear given.
Aug 2, 2021	EDPB	Discloses a positive official opinion on IIJ BCR-C/P.
Aug 5, 2021	LDI-NRW	IIJ BCR-C/P approved.

Table 1: IIJ's BCR Approval Process

The IIJ Group is also aware that it cannot rest on its laurels just because it has obtained BCR approval under the EU GDPR. We realize that, at the very least, we will also need to comply with the UK's requirements now that it is out of the EU, and that we will need to adhere with other personal data protection initiatives, including those in the US. The situation now is such that we could never keep up with all of the initiatives out there, but broad frameworks such as APEC CBPR, at least, are something we believe we should also look at complying with.

IIJ is also not immune to personal information breaches and other incidents, so I think we could be exposed to criticism along the lines of "look who's talking." But given the crucial role of personal information protection and other aspects of information security in our modern information society, I think there is a clear duty to address the issues in front of us. Organizations need to respond to the personal data protection and other practices found in different cultures while also enhancing their own information security capabilities.

Personal information protections and other information security measures are not something that can be thrown together overnight. They become meaningful only once they are established as part of organizational culture and everyday work practices. At the risk of being repetitive, obtaining BCR approval for IIJ was never our end goal. Our bigger objective was to use it as an opportunity to further bolster our internal controls around information security and privacy protection. We are proud to have taken a major, if not definitive, step toward that goal.

I think the IIJ Group's corporate mission can be described as the single-minded pursuit of business with the Internet at its core. Companies that specialize exclusively in global Internet infrastructure seem to have become rare these days. So, as a company that has helped drive the spread of the Internet, I think we have a responsibility to hold on to our Internet-centric perspective and continue making contributions as an infrastructure company for the information society age.

To be honest, IIJ BCR was a slog, and I honestly wouldn't recommend it to anyone. But I can say emphatically that I'm glad we did it because I believe it will be one key pillar supporting the type of Internet we envision.

Looking ahead, we will continue to strive for a safe and secure Internet.

Takamichi Miyoshi

Senior Fellow, DPO (Data Protection Officer), IIJ. Mr. Miyoshi joined Internet Initiative Planning Inc. (now Internet Initiative Japan Inc.) in April 1993. He worked on launching Internet services and the operation of service equipment/facilities. He later served in services development and strategic planning, going on to participate in numerous study groups set up by the Ministry of Internal Affairs and Communications and other agencies as Managing Director of IIJ. He has been in his current position since June 2015.

February 2022 Vol.53

IIJ
Internet Initiative Japan

About Internet Initiative Japan Inc. (IIJ)

IIJ was established in 1992, mainly by a group of engineers who had been involved in research and development activities related to the Internet, under the concept of promoting the widespread use of the Internet in Japan.
IIJ currently operates one of the largest Internet backbones in Japan, manages Internet infrastructures, and provides comprehensive high-quality system environments (including Internet access, systems integration, and outsourcing services, etc.) to high-end business users including the government and other public offices and financial institutions.
In addition, IIJ actively shares knowledge accumulated through service development and Internet backbone operation, and is making efforts to expand the Internet used as a social infrastructure.

IIJ-MKTG020-0051

Internet Initiative Japan Inc.

Address: Iidabashi Grand Bloom, 2-10-2 Fujimi, Chiyoda-ku,
Tokyo 102-0071, Japan
Email: info@iij.ad.jp URL: https://www.iij.ad.jp/en/