IIR

Internet
Infrastructure
Review

Feb.2023
Vol.57

Periodic Observation Report

Internet Trends as Seen from IIJ Infrastructure—2022

Focused Research (1)

VX—IIJ's New Backbone Network

Focused Research (2)

illumino—IIJ's Internal Data Analytics Platform

IIJ
Internet Initiative Japan

Internet Infrastructure Review

February 2023 Vol.57

© Internet Initiative Japan Inc.

Executive Summary

IIR Vol. 57 is our last issue covering 2022. One major topic in Japan's ICT industry over the year was the huge impact that the suspension of ICT companies' services had on activity throughout our society. Television news programs covered the suspension of services not only by traditional telecommunications carriers but also by cloud service providers. And this was not limited to Japan. In South Korea, as well, it was reported that disruptions to smartphone apps with many users affected social life across a wide range of areas.

It goes without saying that ICT companies should be working to improve service quality. In the case of Japan's mobile communications, vigorous discussion about roaming between mobile carriers during emergencies is ongoing. The idea here is that not only should users take steps to protect their own interests by, for instance, arranging alternative means of communication, but service providers should also back each other up to ensure continuity of service in the event of disruptions.

In the mobile communications industry, progress is being made globally on the infrastructure sharing front in pursuit of efficiency, and we are seeing gradual change versus the traditional approach of competing to expand coverage areas by building infrastructure. And when it comes to emergency roaming, too, this shift from infrastructure building being seen as an area of competition to it being an area of cooperation also looks to be in play.

With a new order and new ideas being discussed and worked out in the ICT industry, including with respect to international cloud services, we are keeping a close eye on what we believe to be signs that the fundamental mobile communications business frameworks are being probed from different perspectives.

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 2022 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, and IPv6 and mobile traffic. The Internet is constantly changing, and we believe that considerable insight can be gained from continuously analyzing, on both an absolute and proportional basis, data that is easy to overlook in the midst of such change. This is the rationale behind our periodic observation reports. The results this year are again quite intriguing, and I encourage you to take a look.

The focused research report in Chapter 2 introduces VX, IIJ's new backbone network. Released in 2022, VX was built for use with new services and thus implements features such as SDN, automation, and API linkages. The report explains the background to VX's development and its implementation, and also goes over the history of the backbone network so far. It also provides some insight into our thinking about the IIJ backbone network, which has been at the core of IIJ's business since its founding.

Chapter 3 is our focused research report on illumino, the analytics platform we built within IIJ. The services and systems we operate generate massive amounts of data. Collecting and analyzing that data is something that creates meaningful value for both us and our customers. The report discusses the possibilities that the collection and analysis of data open up and what sort of things we are trying to achieve, along with an example of this in action.

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's MVNO Council, and in June 2021, he became a vice-chairman of the association.

© Internet Initiative Japan Inc.

Internet Trends as Seen from IIJ Infrastructure —2022

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. In particular, we analyze changes in trends from the perspective of BGP routes, DNS query analysis, IPv6, and mobile.

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).

The annual increase in the number of routes returned to over 50,000, with the total number of routes surpassing 900,000. Yet we are observing a downtrend in this route growth off of the 2018 peak (Figure 1), so what the following year's figures will bring is something that is already on our mind. The total number of unique IPv4 addresses increased by a bit less than 32 million (roughly double vs. the year before last), but when the impact of routes advertised by AS749, which account for last year's large increase is excluded, the number of routes actually looks to have fallen by around 1.16 million.

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).

While the total number of routes surpassed 150,000, the size of the increase was only a bit over 50% of the previous year's (around 23,000 routes). The increase in the number

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. 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
Sep. 2022	16	13	39	101	298	592	1208	2064	13502	8292	13909	25051	43972	52203	109071	96909	536520	903760

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. 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
Sep. 2022	298	4247	895	21926	15147	12509	4108	13840	6994	73244	153208

© Internet Initiative Japan Inc.

of unique /64 blocks, meanwhile, was around 3.4 times larger than in the previous year (71.46 billion blocks). This is likely due to the large increase in the number of short-prefix routes (/20 – /31), from which we surmise that the rollout of IPv6 by large network organizations (mobile communication carriers etc.) has progressed. Note that routes for which there is no information on shorter prefixes, which contribute to the additional number of unique blocks, accounted for 45.2% of the increase.

Lastly, let's also look at IPv4/IPv6 full-route Origin AS figures (Table 4). In the past year, an additional 4094 32-bit-only AS numbers were allocated to APNIC and 1024 to LACNIC. Both the decrease in 16-bit Origin Autonomous System Numbers (ASNs) and the increase in 32-bit-only Origin ASNs were smaller than in the previous year. The latter was less than 40% of the previous year's increase, and as a result, 32-bit-only ASNs only came to 49.0% as a proportion of all Origin ASNs. IPv6-enabled ASNs, which advertise IPv6 routes, also increased, but with the growth rate slipping below 10% for the first time in the last 10 years. While the changes were relatively small this past year, whether this trend continues or whether it turns out to be a temporary phenomenon due to the contraction in economic activity amid the Covid pandemic is something that will bear watching for in the coming year's figures.

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. 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
Sep. 2022	3,068,374,784	532,578,391,219

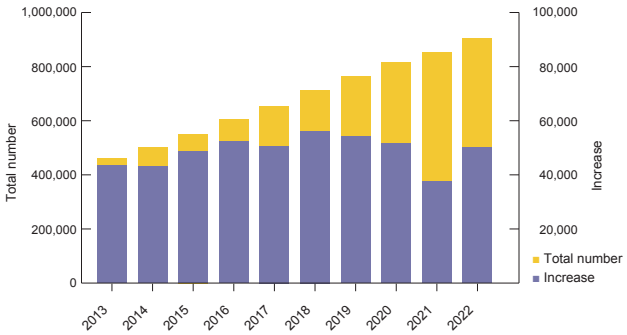

Figure 1: Total Number of Full IPv4 Routes and Annual Increases

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. 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%)
Sep. 2022	11613	28398	369	40380	(29.7%)	10816	22211	5764	38791	(42.7%)

© Internet Initiative Japan Inc.

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 5, 2022.

The full resolver provides a name resolution function that replies to DNS queries from user devices. 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 based on a policy that differs from the 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 0.13 queries/sec per source IP address at around 4:25 a.m., and a peak of about 0.34 queries/sec per source IP address at around 10:00 p.m. The overnight trough is only up 0.01pt vs. 2021, not a huge change, whereas the evening peak has grown by 0.04pt. The growth rates look to have slowed a bit vs. 2021, but the uptrend is ongoing. The breakdown shows that IPv4 accounted for around 59% of queries and IPv6 for around 41%, pretty much the same pattern as in 2021.

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 6 a.m. and 7

© Internet Initiative Japan Inc.

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. Mirroring the pattern also observed in 2021, 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 60% are A record queries and 20% AAAA record queries (Figure 2). With IPv6-based queries, meanwhile, AAAA record queries account for a higher share of the total, with around 40% being A record and 36% being AAAA record queries (Figure 3). Compared with the previous year, we observe 4-percentage-point drops in A record queries for both IPv4 and IPv6.

HTTPS-type records, which we started to see in 2020, accounted for some 15% of IPv4 and 21% of IPv6 queries, marking steady increases of 4 percentage points for IPv4 and 3 percentage points for IPv6. Meanwhile, SRV record queries have fallen as a percent of total for both IPv4 and IPv6, and thus we now group these into the “other” category.

Also in the IPv6 space, we are seeing an increase in new SVCB record queries, although they still only account for a meager 0.12% of total. This may be attributable to Discovery of Designated Resolvers (DDR), a newly proposed protocol for allowing clients to detect encryption-capable full resolvers. In this proposed protocol, a client configured with an unencrypted resolver first queries the _dns.resolver.arpa. SVBC record. When replying to this query, the resolver can include the necessary information to inform the client of encrypted resolvers that support DNS-over-HTTPS (DoH), DNS-over-TLS (DoT), or DNS-over-QUIC (DoQ). Many IPv6-capable clients are newer implementations subject to software updates, and we surmise that this may explain the increase in these queries using the new specification.

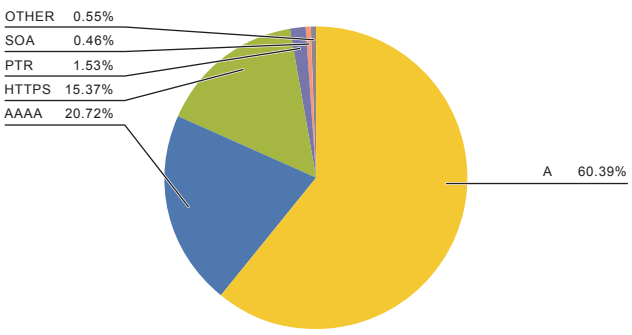

Figure 2: IPv4-based Queries from Clients

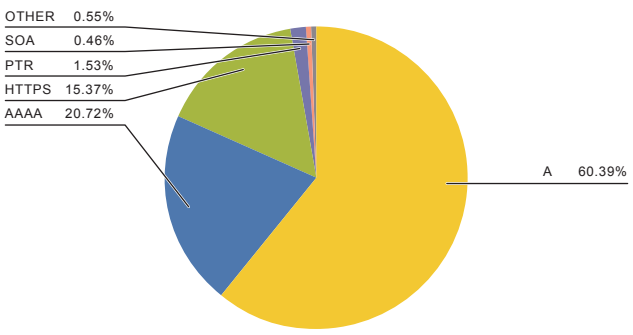

Figure 3: IPv6-based Queries from Clients

© Internet Initiative Japan Inc.

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. Also in this edition, for the first time in three years since we last covered the topic in IIR Vol. 45 (https://www.iij.ad.jp/en/dev/iir/045.html), we go over the state of IPv6 connections according to differences in mobile device OS.

Traffic

Figure 4 shows traffic measured using IIJ backbone routers at core POPs (points of presence—3 in Tokyo, 2 in Osaka, 2 in Nagoya). The data cover the nine months from January 1 to September 30, 2022.

IPv6 and IPv4 traffic was generally range-bound through 2022. This is fairly odd given that IPv4 had so far been growing at a rate of a few percent and IPv6 at a rate in the 10–20% range. Looking back over the past few years, we note a slight stalling in 2020 due to Covid followed by a large rebound in 2021, and this may be why there is no notable trend this year.

Like last year, Figure 5 graphs traffic indexed to 1 as of January 4, 2022, the first business day of the year. Both

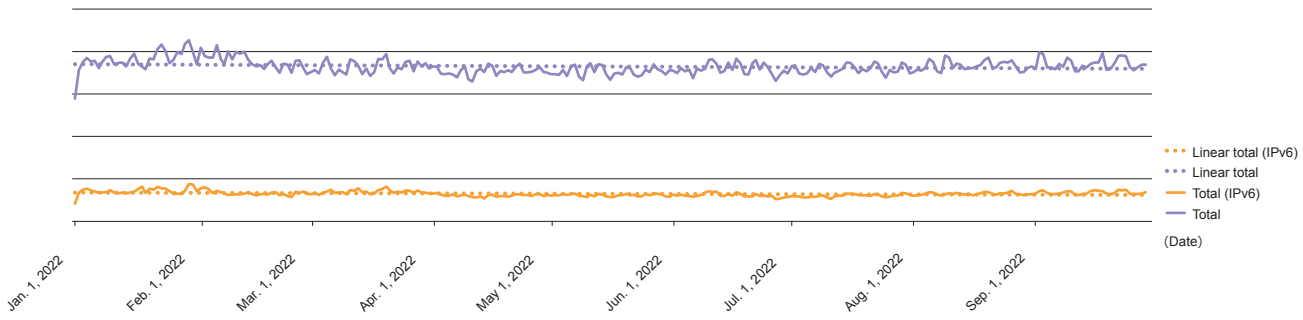

Figure 4: Traffic Measured on Backbone Routes at IIJ's Core POPs

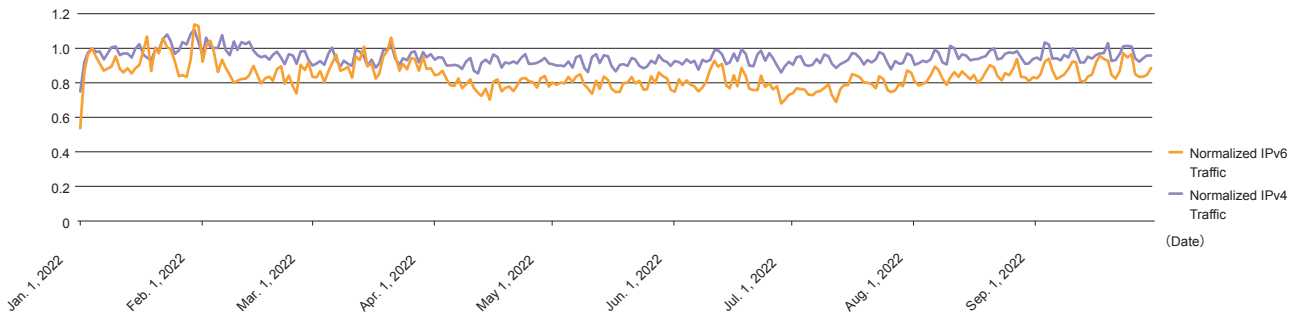

Figure 5: Traffic Indexed to 1 as of January 4

© Internet Initiative Japan Inc.

IPv4 and IPv6 look to have declined slightly rather than staying completely flat.

Next, Figure 6 shows IPv6 as a proportion of total traffic. While it exceeded 22% at the start of the year, it generally moved in a range of 16–20%, with the nine-month average being 17.8%.

Table 5 tracks the IPv6 ratio over the past five years.

Traffic Source Organization (BGP AS)

Next, Figures 7 and 8 show the top annual average IPv6 and IPv4 traffic source organizations (BGP AS Number) for January 1 – September 30, 2022.

In our previous edition of this report in IIR Vol. 53 (https://www.iij.ad.jp/en/dev/iir/053.html), we reported that Company A, a major Japanese content provider, ranked second in IPv6 traffic. In 2022, though, this provider came in at No. 1 with a share of 8.8%. At No. 2 with 7.9% was Company B, which had held the top spot until 2021, and at No. 3 with 3.6%, like in 2021, was Company C, a major US CDN operator. Company G, a major Japanese content provider, is also working its way up with a share of 1.7% + 0.6% (has multiple ASs for each business), and it is evident that Japanese content providers are gradually rolling out IPv6 support.

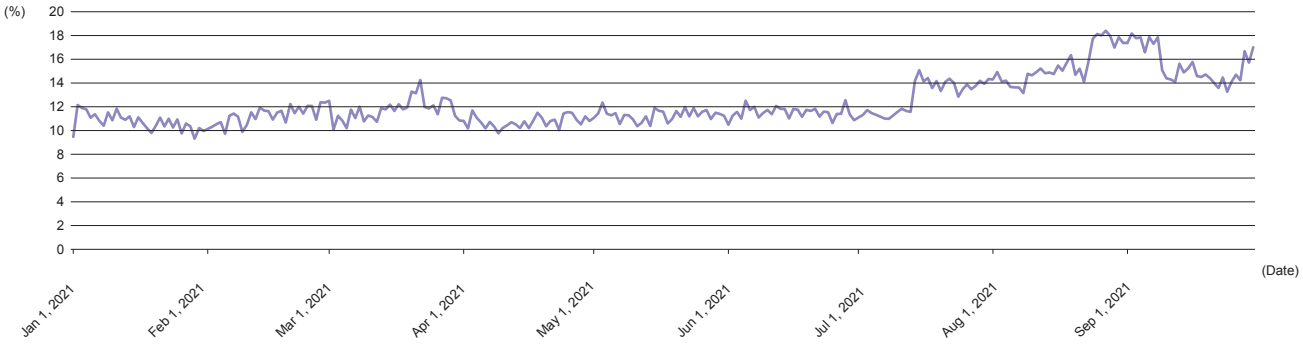

Figure 6: IPv6 as a Proportion of Total Traffic

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%

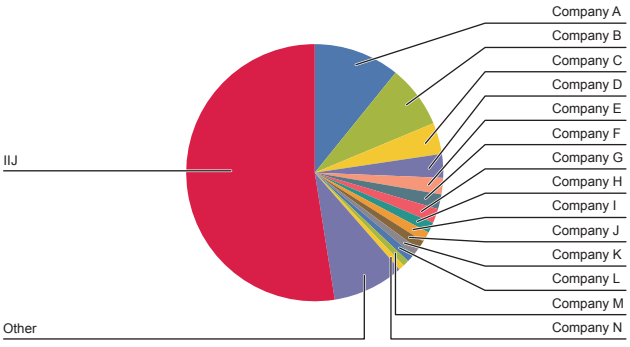

Figure 7: Annual Average IPv6 Traffic by Source Organization (BGP AS Number)

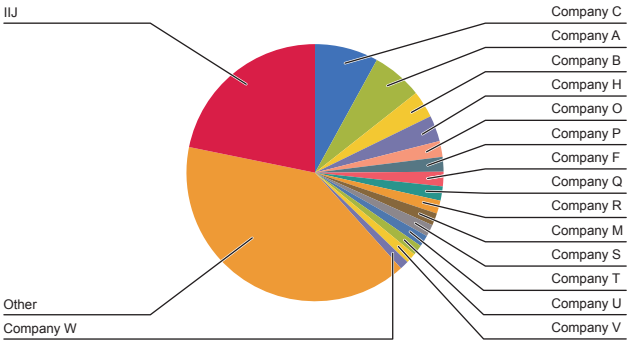

Figure 8: Annual Average IPv4 Traffic by Source Organization (BGP AS Number)

© Internet Initiative Japan Inc.

Protocols Used

Figure 9 plots IPv6 traffic according to protocol number (Next Header) and source port number, and Figure 10 plots IPv4 traffic according to protocol number and source port number (for the week of Monday, October 3 – Sunday, October 9, 2022).

In the IPv6 space, ESP (IPSec) fell from 3rd place in 2021 to 5th place while UDP4500 (NAT Traversal IPSec), in a figurative role reversal, moved up from 5th to 3rd place. NAT is basically not used on IPv6, but perhaps the use of NAT-T ports has to do with standardizing of implementations. IPv6 traffic has also grown during the daytime on Saturdays and Sundays relative to weekdays, but ESP and UDP4500 traffic is down considerably, so one can imagine they are being used mainly for remote work and the like.

The IPv4 trends look mostly unchanged from 2021. Interestingly, IPv4 traffic seems to be falling a little more on weekends than on weekdays, while IPv6 traffic seems to be increasing. This could perhaps mean that IPv6 usage rates are higher in the home than on corporate networks. This is only speculation as we have no definite evidence at this point, but it is something we hope to investigate if the opportunity presents.

IPv6 Across Different Mobile Device OSs

In IIR Vol. 45 in 2019 (https://www.iij.ad.jp/en/dev/iir/045.html), we presented the results of our investigation into whether there were any differences across mobile OSs in terms of whether IPv6 was enabled or disabled on personal mobile service (IIJmio Mobile Service) connections. At the time, IPv6 was enabled on 48% and disabled on 52% of all

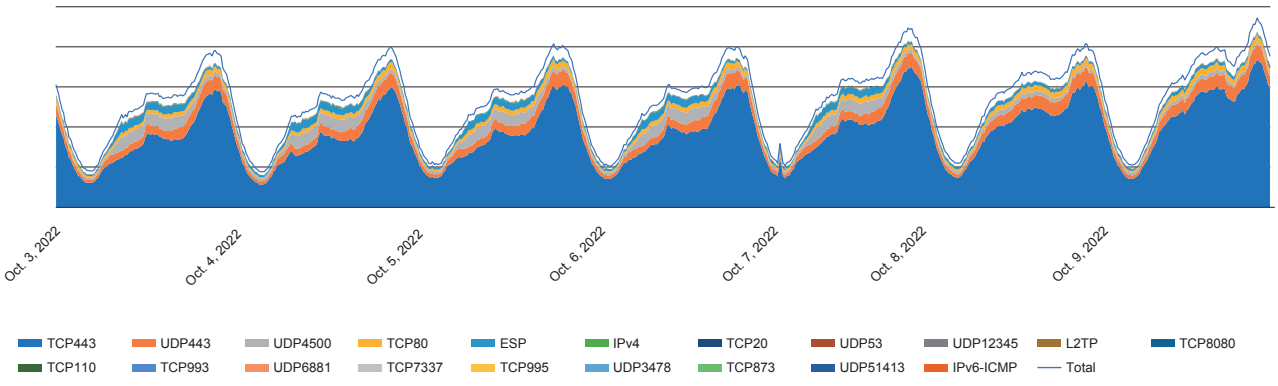

Figure 9: Breakdown of IPv6 Traffic by Source Port Number

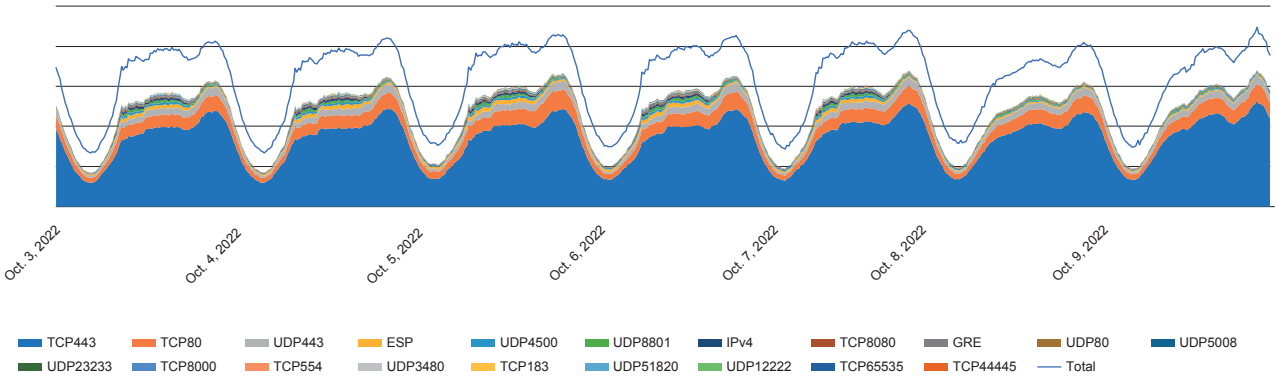

Figure 10: Breakdown of IPv4 Traffic by Source Port Number

© Internet Initiative Japan Inc.

IIJmio Mobile connections, so less than half of connections were over IPv6.

On this occasion, we look at connection rates and device OSs based on data from around 10:30 a.m. on Monday, October 17, 2022.

First, as Figure 11 shows, IPv6 is now enabled on over half of connections (56.3%). Also, we have not included a graph here, but the traffic ratio was 7 (IPv4) : 3 (IPv6) (around 6:00 p.m. on weekdays, when traffic is relatively heavy). Clearly, it seems that IPv6 has seen substantial growth in terms of both number of connections and traffic.

Next, we look at the IPv6 connections across mobile OSs. To do this, the pie chart in Figure 12 breaks mobile devices connected to the IIJmio Mobile Service into three groups—Apple iOS (iPhones and iPads), Android, and other (mobile routers and dongles etc.)—based on part of the IMEI (TAC: first 8 digits) matched against the GSMA database.

IPv6 was enabled on a fairly high 85.7% of iOS connections. That said, IPv6 was enabled on 90.8% of connections when we reviewed the data three years ago, so the proportion has declined.

IPv6 was enabled on 21.7% of Android connections. This is a 7.6-point increase from our 14.08% IPv6 reading for Android three years ago. Yet there still remains a huge difference in the IPv6-enabled rate relative to iOS.

For other devices (Wi-Fi routers, USB dongles, IoT devices, etc.), IPv6 was enabled on 25% of connections. It is surprising that IPv6 is enabled on more of these devices than on Android, and this possibly indicates that IPv6-capable Wi-Fi routers are also on the rise. That said, almost all of the devices connected were smartphones and tablets, so these other devices only have a minimal presence.

Summary

We have examined traffic on the IIJ backbone core, source ASNs, and main protocols used. Although traffic volumes were range-bound or in a slight decline from the beginning of the year, IPv6 usage rates increased vs. a year earlier, reaching a six-year high. Data on source ASNs showed that the IPv6 traffic of Japanese content providers is growing. There were no major changes with the main protocols, and for both IPv6 and IPv4, web-based protocols remained at the forefront, followed by VPN protocols.

We have also looked at IPv6 connections across mobile device OSs. On mobile services, over half of all terminals were IPv6-enabled, and around 30% of traffic was IPv6. By OS, IPv6 was enabled on over 80% of Apple iOS devices, while IPv6 was disabled on 80% of Android devices, mirroring the situation three years ago, but the Android IPv6 connection rate had increased.

We will continue to watch the IPv6 situation from a range of angles and provide updates as new developments come to light.

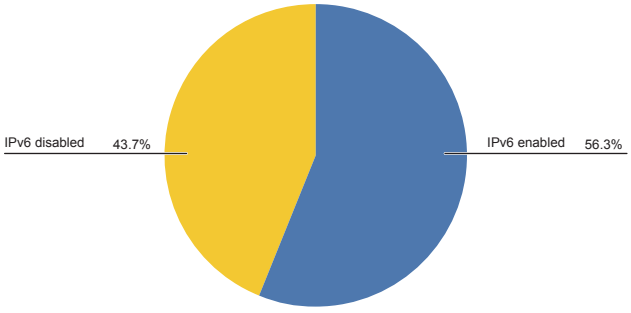

Figure 11: Proportion of Connections with IPv6 is Enabled

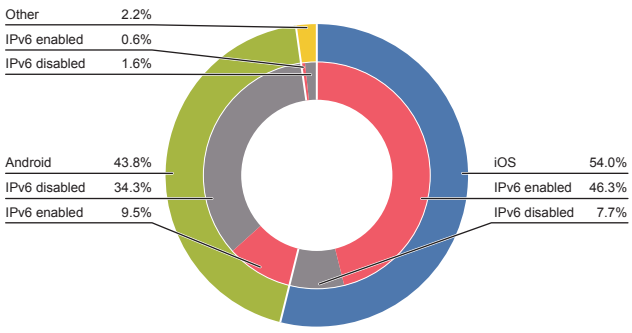

Figure 12: OSs on IPv6-enabled Devices

© Internet Initiative Japan Inc.

Topic 4

Mobile 3G, LTE

Mobile traffic patterns have been affected by the Covid pandemic over the past few years. Here, we summarize what's happened with traffic in the past year, based on observations covering October 1, 2021 – September 30, 2022.

Firstly, NTT Docomo will terminate 3G communication services at the end of March 2026, so we report on the current 3G traffic situation.

3G traffic as a percent of total (Figure 13) is as follows. On consumer services, 3G communications are virtually non-existent, accounting for only around 0.05% of total traffic. In business services, 3G averages 6.4% of total. Looking at the trends, it remained in a very moderate decline up until April 2022, but that downtrend looks to have accelerated since May 2022.

Next, we look at traffic and session counts on business services. Here, we graph traffic volume (Figure 14) and session counts (Figure 15) for business services indexed to October 1, 2021.

Looking at traffic volumes, we see that LTE traffic volume remained in a gradual uptrend throughout the year, with that uptrend appearing to have accelerated just slightly from June 2022. And as mentioned above, 3G traffic volume was in a slight decline up until April 2022, with the downtrend then accelerating from May 2022.

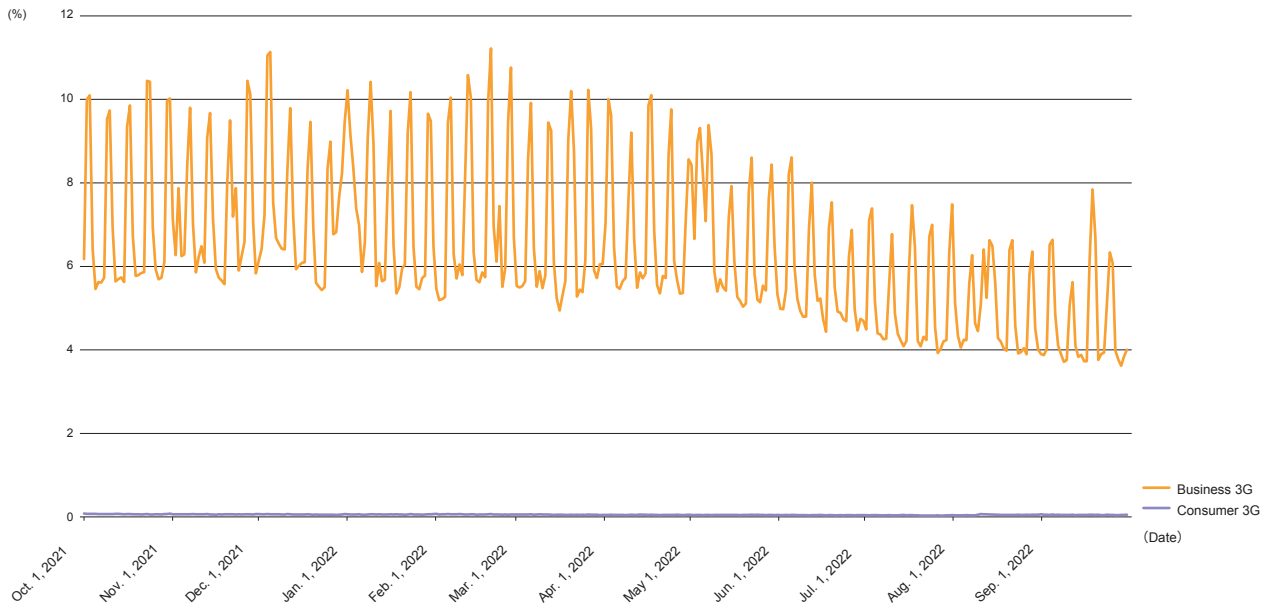

Figure 13:3G Communications as a Proportion of Total Traffic

© Internet Initiative Japan Inc.

Looking at session counts, we see that, similar to traffic volume, LTE session count was in a gradual uptrend throughout the year, with that uptrend appearing to have accelerated just slightly from July 2022. The 3G session count remained roughly in line with the base date of October 1, 2021 up until December 2021, with an intermittent downtrend appearing after 2022 got underway.

On business services, traffic trends are affected by the progress of customers' plans to migrate from 3G to LTE. As someone in charge of mobile services equipment, I am pleased to see 3G decline, given that it is on the way, and LTE increase, so we will be keeping tabs on the decline in 3G communications as we continue to provide stable infrastructure.

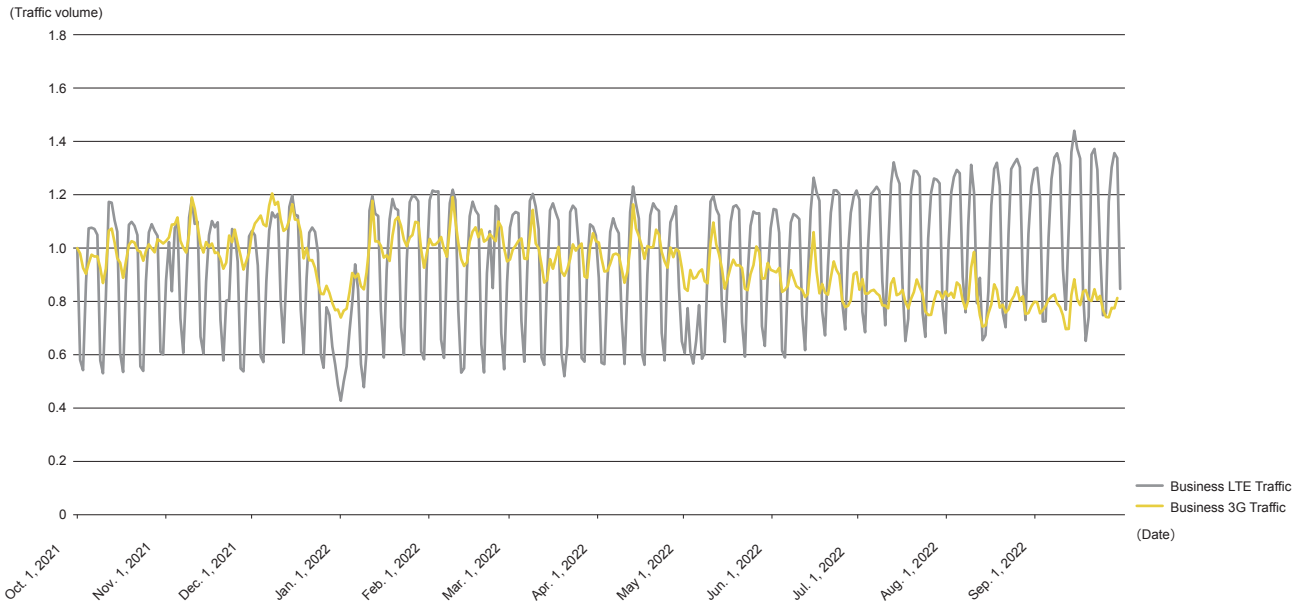

Figure 14:Traffic Volume on Business Services

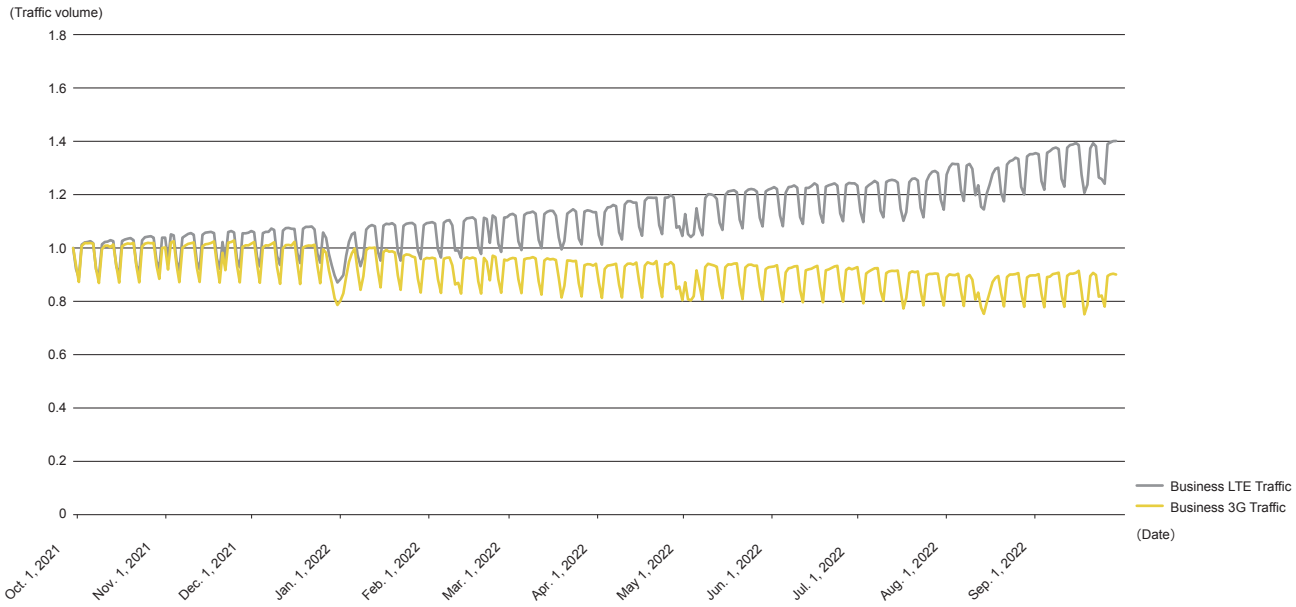

Figure 15:Session Counts on Business Services

© Internet Initiative Japan Inc.

Next, we look at traffic and session counts on consumer services. Here, we graph traffic volume (Figure 16) and session counts (Figure 17) for consumer services indexed to October 1, 2021.

In terms of traffic volumes, LTE remained largely range-bound up until late February 2022 amid Covid quasi-state of emergency measures implemented by the Tokyo Metropolitan Government, but it increased steadily from

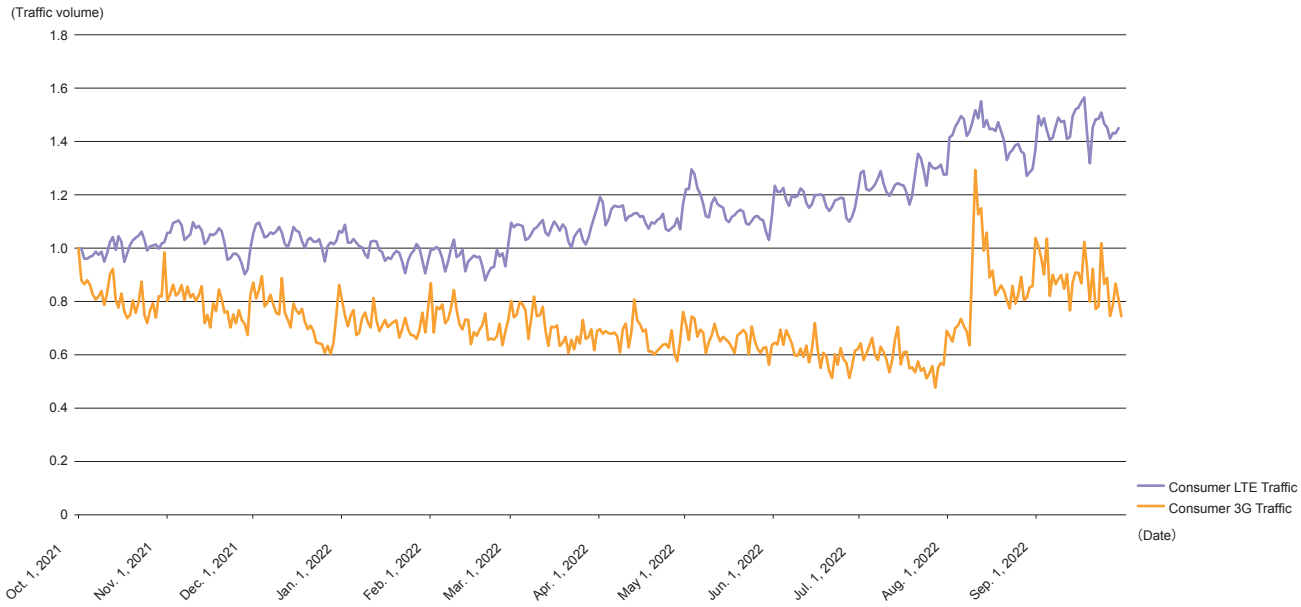

Figure 16: Traffic Volume on Consumer Services

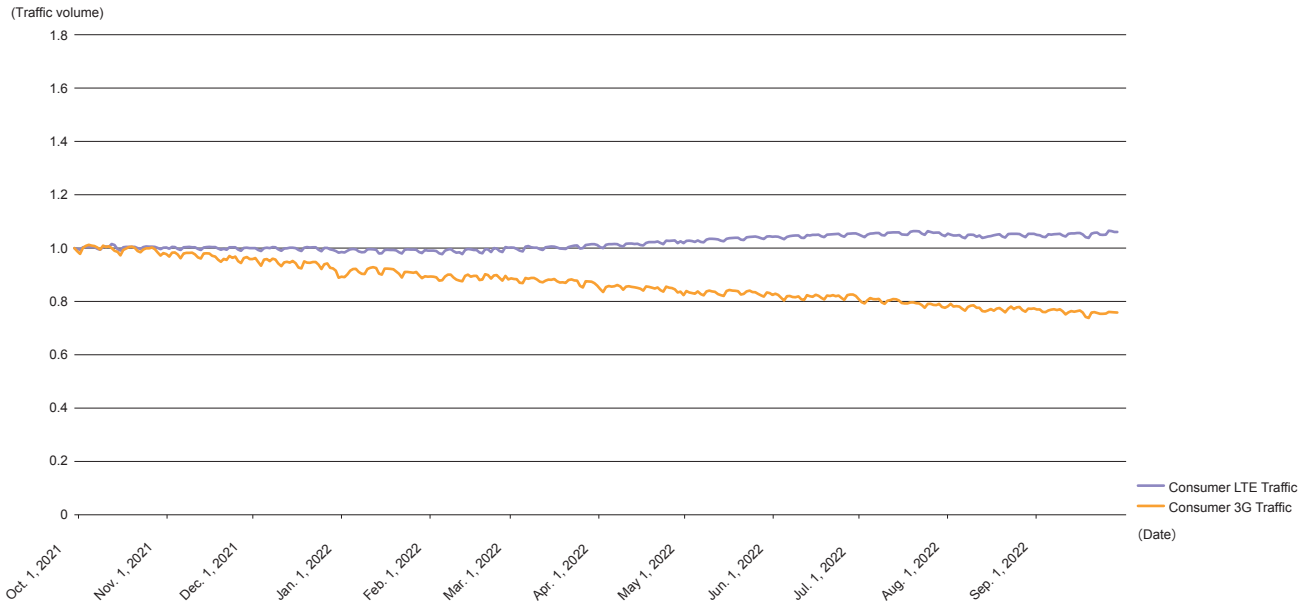

Figure 17: Domestic Interconnection Traffic

© Internet Initiative Japan Inc.

March onward, reaching roughly 1.4x the index date's level. Traffic volume increased substantially from early August in particular, but this is attributable to enhancements in infrastructure capacity. 3G traffic also had been in a steady decline but surged sharply in early August, similar to LTE traffic. This probably reflects infrastructure capacity enhancements, as with LTE. While it does look like there was a large effect on 3G traffic, as explained earlier, 3G communications are virtually non-existent relative to LTE as a proportion of total on consumer services, so even a small effect will appear as though there has been a large impact.

Looking at the session counts, meanwhile, the LTE session count remained largely range-bound throughout the year with a modest increase. And the 3G session count stayed in a downtrend throughout the year.

LTE accounts for almost all of the communications on consumer services, and although session count remained unchanged, traffic increased to around 1.4x the index date's level over the year. In other words, traffic volume per session is simply increasing. This no doubt means that the more you can accomplish on a smartphone, the more traffic per session will tend to increase, and while this may raise all sorts of difficult issues on the infrastructure side of things, we hope to continue coming up with solutions going forward.

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
Mobile Technology Department, Infrastructure Engineering Division, IIJ

4.State of the Mobile Industry and Traffic Trends
Tsuyoshi Saito
Deputy General Manager, Mobile Technology Department, Infrastructure Engineering Division, IIJ

© Internet Initiative Japan Inc.

VX—IIJ's New Backbone Network

2.1 Introduction

In June 2022, IIJ began offering a new network service called IIJ Private Backbone Service/Smart HUB (the "SHB service"), which provides flexible high-capacity connections with cloud services. We built and released a new backbone network, which we (internally) call VX (Virtualization eXchange), to provide the network infrastructure for the SHB service. This article gives an overview of this new VX backbone network and the background to its creation, and details, from various angles, how it differs from previous backbone networks.

2.2 History of the IIJ Backbone

I would like to start by unraveling the history of IIJ's backbone network, which I don't think has been told in much detail to date.

VX is the fourth-generation of network infrastructure within the context of IIJ's backbone network. Let's take a brief look at each generation. The first generation is the Layer 3 IP network ("BB") that has been around since IIJ's early days. BB started with a 192kbps circuit back when IIJ was established, and it has since expanded beyond Japan to now encircle the globe. While originally using 192kbps lines, as of 2022 it mainly runs on 100Gbps broadband, and we are now looking at deploying next-generation 400Gbps. It provides Internet infrastructure for a range of services, starting with IIJ's Internet connectivity services, and continues to expand in that regard.

The first-generation backbone network that is BB reached a turning point in the 2010s. Until then, BB followed a policy of using the same physical and logical topologies, and the network was designed to carry traffic as efficiently and stably as possible over its routers and circuits. Back then, BB took on a square configuration with two core routers each in Tokyo and Osaka, and the leaf nodes in eastern Japan were connected in a V-shaped configuration. Although this network topology may seem efficient at first glance, the bandwidth at the busiest point between Tokyo and Osaka had to be kept constantly at 50% or lower. Also, in the event of a failure or maintenance on one part of the system, all traffic had to be diverted to the other part of the system, and the operational burden and costs involved in doing this were high. So to maximize fault tolerance and traffic-balancing efficiency, we created a mesh network, called a backbone fabric (BF), between the core routers using carrier circuits and our own WDM equipment, making it possible to connect the core routers in an N+1 configuration. At the time, the BF was made up of six routers in

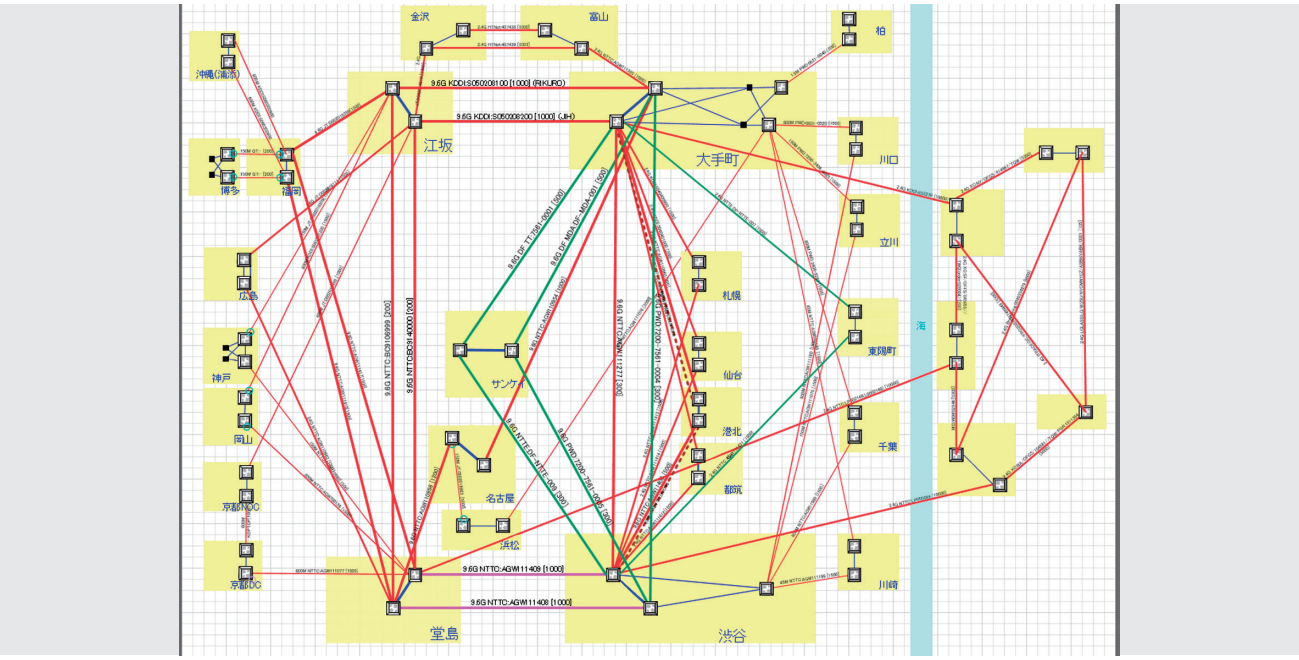

Figure 1: IP Backbone Map Circa 2006 (when Osaka–Tokyo was in a Square Topology)

© Internet Initiative Japan Inc.

Tokyo and four in Osaka, and we ultimately expanded it to the US East Coast and West Coast. Using the BF topology drove traffic efficiency on the Tokyo–Osaka leg all the way up and, I think, made it possible to achieve stable communications with minimal impact on traffic relative to how things were with the core routers in a square configuration. We created the BF using 10Gbps Ethernet and 9.6Gbps SONET/SDH circuits, but not being a carrier, IIJ did not have its own carrier network. A disadvantage of not having such a network was the increase in costs posed by circuit usage fees, but, more importantly, the advantages were the ability to freely select circuit operators and to procure from a wider range of circuit path options for the BF topology, which involves using a lot of circuits.

While BF did achieve this ideal, the time came when we started to run up against limits to maintaining this configuration. To make maximal use of BF, the core routers need to be fully meshed with the BF, and a lot of work needed to be done when increasing speed to create the many backbone links. The network topology policy at the time was that the physical topology should match the Layer 3 logical topology. But as we ran up against limits to maintaining this physical topology, we decided to rethink this notion that physical topology should mirror logical topology. This heralded the evolution of the IIJ backbone network into its second generation.

The second generation was a virtual Layer 2 network that we called WARP. The concept with WARP was to network sites in a way that is independent of the physical topology, and to thereby address issues experienced with BF. So, starting with WARP, we began creating logical paths using MPLS label switching technology, something that had not been used on the IIJ backbone network until then. In the BF era, we created Layer 3 load balancing paths along physical circuits, but with WARP, the virtualization of the network between sites meant that we could freely create Layer 2 connections between sites that were not directly physically connected, so we had a greater degree of freedom in terms of topology. WARP facilitated full-mesh connectivity between network nodes physically configured into a BF. As of 2022, we continue to maintain and expand our WARP-based network between real-world sites, and the core BB routers—those in Japan + key routers in the US—are fully meshed in the form of a virtual Layer 2 network. And with OSPF on network Layer 3, we have achieved optimal logical path topology, enabling precise traffic engineering.

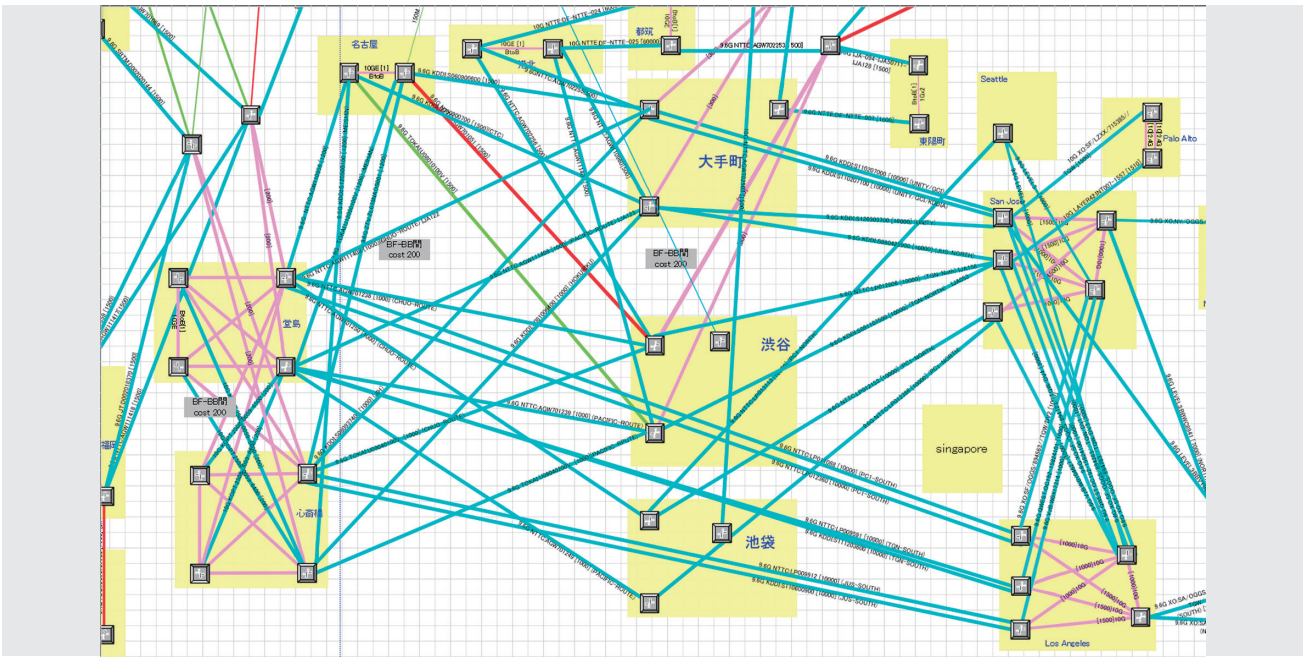

Figure 2: IP Backbone Map Circa 2013 (Backbone Fabric Era)

© Internet Initiative Japan Inc.

Our focus with the network up to this point had been on the efficient and stable transmission of Internet traffic generated between IIJ's sites, but with the third-generation backbone network we turned our attention to links between separate service infrastructure. And in the mid-2010s, we began providing our third-generation backbone network service infrastructure, which we called MATRIX. The concept with MATRIX was to provide a wide-area private network connecting multiple points, and to facilitate interconnections with separate service infrastructure networks, which had generally been independent until that point. WARP was a network for providing virtual Layer 2 connections, but MATRIX was a Layer 3 VPN infrastructure connecting different Layer 3 networks via private networks. Before the advent of MATRIX, connections between different Layer 3 networks (aside from those established via private networks set up for that purpose) required each set of service infrastructure to have a global IP address, and network connections between service infrastructure were generally made via the first-generation IP backbone. The issue with setting up a separate closed network is that this would always involve a bit of effort, in terms of preparing multiple WARP circuits and routing traffic through private-edge routers for an Internet VPN. So setting up an independent Layer 3 network as a backbone meant that it was easy for service infrastructure administrators to establish the necessary interconnections between networks without having to worry about private network issues. It may seem quite obvious to say this now that we are well into the heyday of the cloud, but there is a strong need for different private networks and private networks that do not go through the Internet. The flexibility in network connections between different sets of service infrastructure facilitated by MATRIX has bolstered the network linkages between those service infrastructures and made IIJ's services even more flexible. To this day, MATRIX is helping to facilitate the expansion of IIJ's GIO cloud services and services providing private connections with a range of public clouds, facilitating high-quality private network services that benefit many customers.

2.3 The VX Concept and the Introduction of VX Controllers

So far in this discussion, we've spent a bit of time looking back through the IIJ backbone network's generations, seeing how it has evolved along with the needs of the time and to address certain issues. Our current efforts to deploy VX are also aimed at resolving issues with the IIJ backbone network and realizing service concepts required by today's ICT services. IIJ's customers also continue to make use of the cloud, and use of the cloud has gained

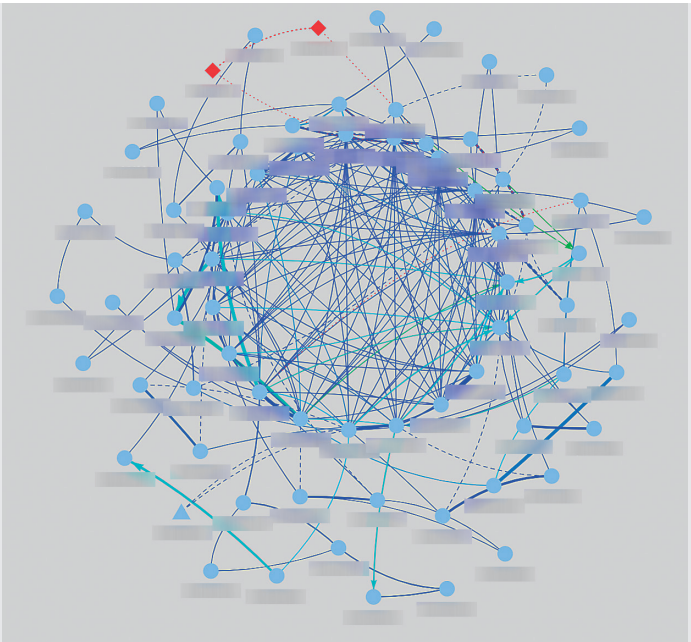

Figure 3: Logical Configuration of the IP Backbone as of 2022
*The meshed part in the center is on the WARP Layer 2 VPN

© Internet Initiative Japan Inc.

further steam amid the spread of Covid since 2020. The volume of business system cloud traffic is set to increase sharply under these circumstances, creating the need for greater network bandwidth. And when it comes to obtaining the necessary bandwidth and cloud resources, users are demanding services that give them the flexibility to use just the amount they need at the time they need it, and this is something IIJ's services need to embody as well. However, while we were able to provide service infrastructure offering network stability and security with MATRIX, we still had issues to address: increases in speed just weren't keeping up with customers' requests for more bandwidth, and operational workloads within IIJ were increasing because our operational style was one of handling the processes involved in delivering services to customers manually. In public clouds these days, the network is abstracted, and users have the freedom to use networks as a cloud feature without having to worry about the physical structure, and they can start up the instances they need from the control panel to make instant use of cloud resources. Hence, it is becoming the norm to provide services in a way that matches the sense of speed users expect. So to ensure that IIJ would be able to provide the services that customers demand, we began looking at the prospects for VX as the fourth generation of our backbone network.

VX is aimed at creating the network infrastructure IIJ needs to quickly and flexibly provide a whole range of services to customers. Our thinking was focused on providing a stable, high-capacity network, as discussed earlier, as well as network infrastructure that could meet the demands of service infrastructure when those services are provided to customers in NFV form. To realize the VX concept, in a first for the IIJ backbone network, we adopted SDN control using a network controller. With recent SDN technologies, you can put everything together from scratch using open source, but when it came to building VX, our chosen approach was to make full use of solutions from vendors that we work closely with on a regular basis. For the initial VX infrastructure, we selected Cisco ACI (Application Centric Infrastructure), an SDN solution for data centers from Cisco Systems. Cisco ACI makes it easy to build a network fabric by using an APIC (Application Policy Infrastructure Controller, an SDN controller) to control network configuration with Nexus series Layer 3 switches as the network nodes. Cisco ACI was originally a solution for making it easy to build IP-Clos networks with a basic spine-leaf topology to serve as server networks within data centers, but we use it for more than just data center purposes at IIJ. With customizations, we also use it for networks connecting POPs between multiple points terminating at end users, NVF server infrastructures, and external public clouds.

The introduction of SDN marked a transition in our backbone network operations from an era in which we mainly operated routers through a command-line interface to an era in which our operations are centered on using controllers, in which we use SDN controllers to, for example, configure network settings and monitor the status of entire networks. The biggest change with the introduction of SDN controllers is that we are now able to use APIs to control networks. The Cisco ACI internal settings are abstracted out to make it easy to use not only for network engineers but for application engineers as well, but even so, it's still fairly daunting for users to deal with directly. So with VX, we used the Cisco ACI API and let users define their own models from the ACI settings based on easy-to-use models, and thus abstracted out the structure to make it look simple. Services are provided in such a way that users can establish connections between the necessary points with only the minimum of VX connection elements and parameters. I think abstracting out and simplifying the structure made it easier to think about API links between VX and the service infrastructures, and easier to make effective use of VX in IIJ's service infrastructure as the core network NFV. VX provides an API interface called VX Controller. Up until now, backbone network operators configured the settings needed to connect customers and service infrastructure, but opening the API makes it possible for VX users to enter settings on demand. We believe that making it possible to use the network infrastructure with as little human intervention

© Internet Initiative Japan Inc.

as possible has shortened availability lead times and made it easier to think about and deploy IIJ services in a more NFV-like manner. We have actually already started using both VX and the API interface, and we use VX as the network infrastructure for providing the SHB service. By linking the back-end API server of the SHB control panel with VX Controller, we have been able to provide an environment that gives customers on-demand network control on the IIJ backbone network and successfully automated the delivery of internal IIJ services.

Figure 4 shows a sample VX Controller GUI screen. Rather than loading in a config, users can create end-to-end connections simply by entering the necessary parameters in the GUI settings panel. An API interface is also provided, which can be used to configure settings just as they would be via the GUI.

If VX is to serve as flexible, high-capacity NFV infrastructure, the network also needs to be easily expanded and extended. For the high-capacity component, we make full use of the high-port-density, high-capacity nature of Cisco Nexus products. Back in the MATRIX era, we also used high-performance routers designed for carrier applications to build the networks, but certain aspects of such high-performance routers got in the way of any meaningful increase in ports—they invariably had low port density and the price-per-port was high. On this point, as a Layer 3 switch, the Nexus is a network device suited specifically to configuring IP fabrics not based on full Internet routes. Deciding on a narrower set of network use cases meant that the requirements were different from those that applied with the high-performance carrier routers previously used on the IIJ backbone network, and this helped us achieve cheaper, high-capacity communications. By narrowing down the functions, as described above, we successfully reduced the initial costs involved in network expansion to an extent. We also need to expand connection points with VX to scale the network and make it easier to connect to VX. Making it easier to expand the network will also no doubt make it easier for IIJ to deploy service infrastructure and easier for more customers to connect to VX. In terms of creating a flexible NFV platform, the expansion of the network also fits in with the idea of moving the exchange of traffic to locations closer to the user in keeping with today's MEC (Multi-access Edge Computing) approach. MEC requires functionality such as routers, switches, and firewalls to be closer to the user, and IIJ can achieve this by deploying VX and an NFV platform as a service close to the customer. This

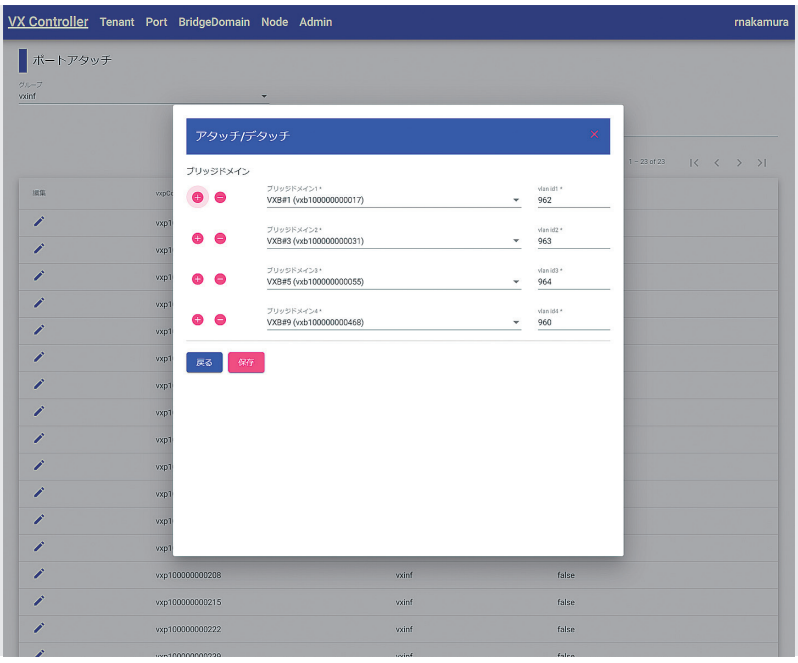

Figure 4: Image of the VX Controller GUI Screen

© Internet Initiative Japan Inc.

makes it easier not only to expand physically but also to logically accommodate the infrastructure for services in the form of an NFV platform. We logically divide each set of service infrastructure based on VX and provide it as a tenanted service. In addition to network divisions, limiting the scope of what can be done on each set of service infrastructure makes it possible to accommodate multiple services virtually. So the impact of one tenant's actions do not spill over to other services, and from the VX user's perspective, anything to do with service infrastructure can be thought of in terms of VX. Control of the API mentioned in the previous section is also limited to the tenant that owns the service, so there is no risk of connections being made to unexpected destinations, which makes it easier for us as the VX provider to pass control over to users.

Let's take a brief overview of VX (Figure 5). It has three layers: a layer for controlling the entire system via the SND controller / VX Controller, as mentioned earlier, a layer that interconnects the data centers, and the spine-leaf fabric network accommodating POPs / NFV platforms / public clouds within the data centers. The nodes are basically configured for redundancy, and thus the system is designed such that a single failure will not affect the provision of services. There are six SDN controllers, including one standby unit, deployed across three data centers. To ensure operational stability, at least three of these units must be running at the same time, and the design ensures that service provision will absolutely not be interrupted even if one of the data centers becomes unavailable. When the number of points or the number of services accommodated increases, we scale out the spine-leaf configuration. It is now easier to deploy services since we are able to expand VX connection points using only the minimum equipment necessary.

2.4 Network Monitoring on VX

Equipment and usage monitoring will also be important as the use of NFV platforms on VX accelerates. Until now, we have generally used ICMP echo for alive monitoring and SNMP for acquiring information and receiving traps when monitoring backbone network devices, but with VX, we have added frameworks to enable the monitoring of network quality from the user's perspective, namely metrics for monitoring network status and tools for

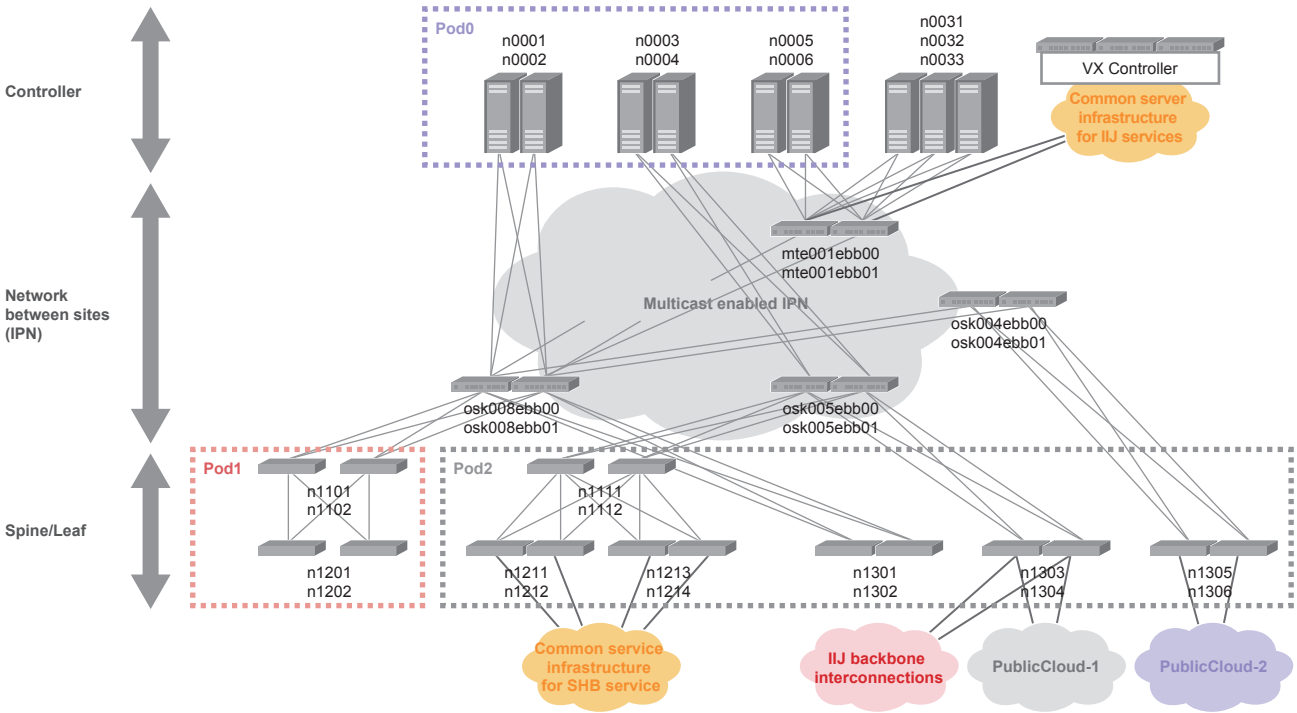

Figure 5: Structural Overview of VX

© Internet Initiative Japan Inc.

effectively visualizing the information. To acquire and store our metrics, we use Prometheus, an open-source tool that is starting to see widespread use of late. Cisco ACI also supports monitoring agents that acquire metrics, so it was easy for us to obtain time-series data. And for data that the monitoring agent is unable to acquire, Prometheus also provides exporters that make it easy to create time series in a specified data format for reading into Prometheus, so a whole range of data can be handled within Prometheus. The metrics collected can be visualized in a dashboard format using Grafana, an open source visualization tool, making it easy to check network health, and we have linked this with our own monitoring system so that alerts can be issued when the system detects monitored values falling below a given threshold. In addition to network device health and errors, the monitoring system can also collect data on service capacity, such as network bandwidth usage and connectable interfaces, to produce visualizations, and we are thus using Prometheus + Grafana to automate capacity checks.

Figure 6 is a sample screen from the VX monitoring dashboard that we use. It gives a comprehensive overview of the status of resources and alerts. By convention, network devices typically appear in green when operating normally, and in orange/red when any anomalies are detected.

We find it difficult to ascertain the status of the network IIJ provides to its customers in the same environment that users experience. While we can monitor service-providing equipment such as the routers and switches that make up the backbone network, there are always some things that monitoring of IIJ's equipment alone will miss. In rare cases, we do unfortunately discover faults only after a customer detects an anomaly and contacts us about it. These are known as silent failures, and they cause disruptions to customers' communications despite no device alarms being generated and no anomalies being present in the logs. Silent failures have long been an issue for us network engineers when it comes to providing services. We have been trying to find ways of detecting silent failures before customers do so that we can swiftly restore service availability. As an NFV platform, we envision VX flexibly interconnecting many different sets of service infrastructure, and we thus expect silent failures to have a significant impact. This is why, with VX, we have introduced mechanisms for monitoring communications status under conditions that are as close as possible to those experienced by the infrastructure users. Every VX service edge is connected to a quality monitoring server, and the servers monitor whether communications via the VX service edge are getting through properly. Since this makes it possible to see communications status from the

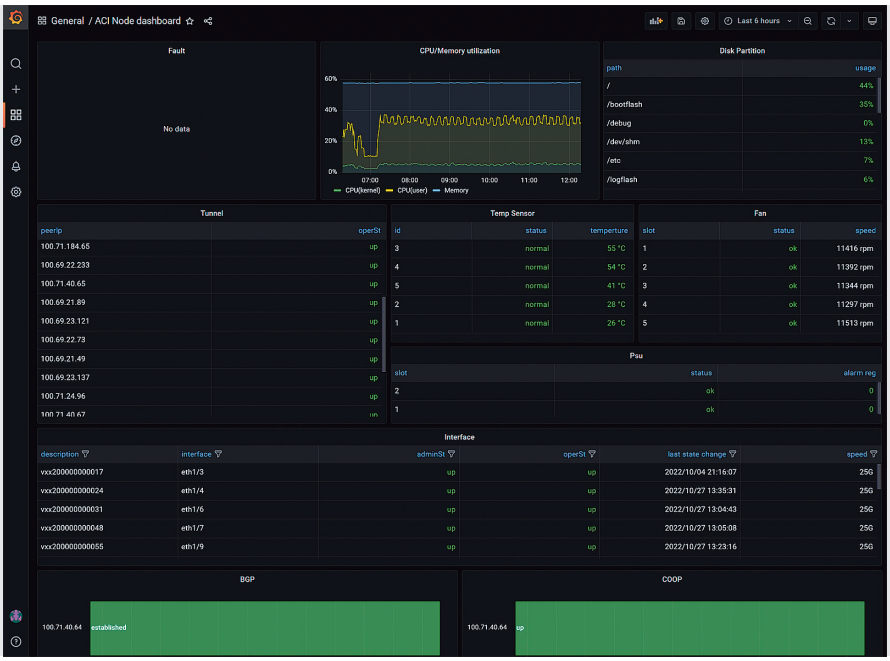

Figure 6: Image of the VX Monitoring Dashboard

© Internet Initiative Japan Inc.

same perspective as users, we can tell if any problems are happening even if the service monitoring system does not detect any alerts. The time and effort required to add on a means for monitoring the status of communications between all service nodes increases as the network grows larger. From the VX initial design phase, we also worked on including a means of monitoring the network from the infrastructure user's perspective, and so we were able to put this into action smoothly to coincide with the VX launch.

Figure 7 shows the screen for monitoring the network from the user's perspective, giving a visualization of the status of communications between nodes over time. The red boxes on the screen correspond to when we actually performed network maintenance, and you can see that there was a partial impact on communications. Also, alerts are sent to the operations center when values fall below set thresholds.

2.5 Conclusion

This article has taken a look back at each generation of IIJ's backbone network and discussed our efforts and concepts for the newly released VX. Last but not least, IIJ has built multiple backbone network platforms across generations one to four, but the release of a new network generation does not mean that we will be discontinuing or merging previous generations. Each network has an optimal role to play and functionality to provide, and our approach aims to use each of the backbone networks synergistically so as to optimize the overall system. All of these networks represent infrastructure that is essential for providing IIJ's services. The recently released VX is not intended to replace previous backbone network generations. Instead, we intend to use this new backbone network to link a whole range of networks, NFV platforms, and cloud services to enable IIJ to deliver services that provide value to its customers.

Figure 7: Screen for Monitoring Network Status between VX Service Edges

Yuichi Yomogida
Part of the team at AS2497. Previously involved in running IX services at JPNAP.
Currently working on the IIJ backbone network and serving as peering coordinator.

© Internet Initiative Japan Inc.

illumino—IIJ's Internal Data Analytics Platform

3.1 Introduction

As the use of IT in business processes continues to expand, many people no doubt feel that information volumes are growing more and more all the time. And commensurate with the volume of information, the number of servers, network devices, and applications managed also increases. This is why the amount of data to be managed is increasing at an accelerating pace beyond the amount we would normally want to deal with.

As the amount of data to be managed increases, devising ways of handling the data efficiently is key. There are many considerations when it comes to storing and managing data, including storage efficiency improvements, data integrity, ease of analysis, increases in analysis speed, and data security. In many projects up till now, only data saving and integrity, or a limited subset required for the task, were used. The data we analyze consists of not only structured data but also a lot of unstructured data such as text logs, which can take a bit of time and effort to analyze. Moreover, application architectures are moving away from large, unwieldy affairs and toward becoming more standardized and decentralized through microservices, and advanced analysis methods are now indispensable.

Being able to use a system capable of managing and making advanced use of data as a common platform would not only make it possible to focus on enhancing the value of services and systems, which is ultimately what we should be doing, but it would also help us to generate new value by using data in ways we were previously unable to. Enter our new data analytics platform, illumino.

This article describes illumino's features, the issues it has so far resolved, and how this was achieved.

3.2 Introducing illuminok

3.2.1 About the illumino Data Analytics Platform

■ Challenges

IIJ has many systems running, the number and scale of which grow every year. So far, the approach to managing the data generated by these systems was to implement the necessary and sufficient measures for doing so on a system-by-system basis. While some projects had evolved with the implementation of advanced data analytics, many services were left operating on a "necessary and sufficient" basis given cost and man-hour issues. Underutilized data here not only represents potential value for those systems but may also embody value for the company as a whole.

Solutions for utilizing data have evolved rapidly in recent years, making it easier than ever to efficiently manage large amounts of data and perform advanced analysis, but the problem with implementing solutions on a system-by-system basis is that it scatters the necessary cost outlays as well as the management and analytical expertise.

■ Solution

To address these issues, we built a common platform that makes it possible to implement data management and analysis easily and at a low cost—namely, the illumino data analytics platform. This solution provides, as a service, the data storage and analysis tools necessary for analyzing data along with system operations and implementation support from dedicated engineers.

This solution is available to anyone running a project within IIJ. Because it is provided as a service, it becomes available immediately after someone applies for it.

© Internet Initiative Japan Inc.

Let's review the benefits of this solution.

- Data storage
  Can store large volumes of data safely at low cost
- Analysis tools
  Advanced tools available
- System operations
  No need to handle this on the user project end
- Implementation support
  Dedicated engineers / data scientists are available to assist

With many of our internal systems now using illumino, the utilization and analysis of data at IIJ has been making forward progress.

3.2.2 What is Data Analytics?

■ ITOA

No doubt many people sense an increasing reliance on IT in their regular activities. With that increasing reliance, a high level of stability in system operations is being demanded. The systems, meanwhile, are increasing in number and the services based on them are becoming larger and more complex. Naturally, amid this ongoing evolution of systems and services, IT operations must also change. To provide the requisite highly stable level of operations, we need to transition from the conventional ITOM (IT Operation Management) paradigm to ITOA (IT Operations Analytics).

Say a particular service experiences a failure. It was often the case with conventional IT operations for the investigation team to examine data managed separately on each system, ask the relevant organizational units to cooperate in the investigation, and repeat the process until the problem was solved. Naturally, as the scale and complexity of services increases, it is often the case that the system on which users are aware of an error, or on which monitoring tools and the like have detected an error, is not the direct cause, and the number of associated systems (i.e., systems suspected of being the cause) also increases.

ITOA, as the name suggests, is a means of solving issues through the analysis of IT operations. Enabling the efficient collection of large amounts of data along with high-speed searching and advanced analytics logic makes it possible to automatically search data (logs) relevant to the error(s) detected when a fault occurs and swiftly track down the root cause. Trying out a range of data analysis also makes it possible to predict potential issues and new areas in need of attention.

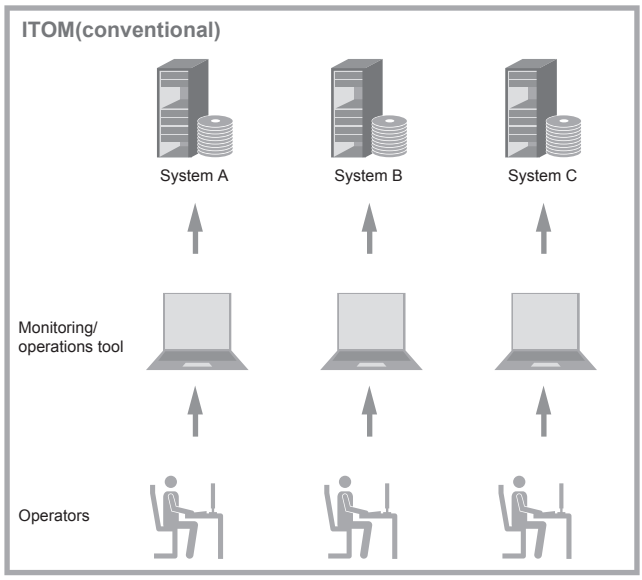

Data is managed on a system-by-system basis

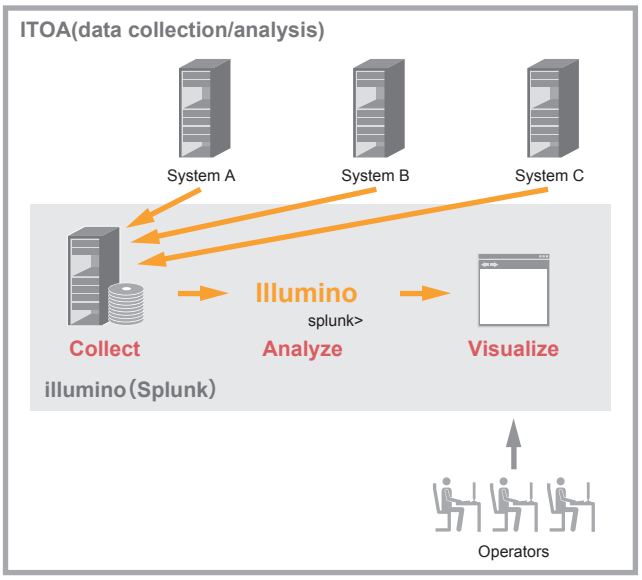

Data is collected and analyzed/visualized
Facilitates integrated operations

Figure 1: How ITOA Works

© Internet Initiative Japan Inc.

Overview of Splunk

Data collection, high-speed searching, and advanced analysis are key to making ITOA work. Of the many ITOA-related solutions out there, IIJ selected Splunk Enterprise from Splunk Inc.

Our reasons for selecting Splunk Enterprise include:

1. Had been used within IIJ before
   We had used it on some projects and thus already had a strong skillset
2. Enterprise grade track record
   Splunk has a track record of being used in large-scale, high-capacity projects worldwide, and we believed it would perform well in the use cases we envisioned for the IIJ common platform
3. Flexible system configuration
   We wanted to build it as an internal IaaS for reasons relating to connectivity with and security of the data input side of the system. But it also offers the flexibility to connect with SaaS offerings such as Splunk Cloud Platform and Splunk Observability.

Data collection

Data comes in all sorts of forms, and in almost all cases it needs to be processed a number of times to achieve the desired result. To achieve this with typical data-using applications, you had to structure the data beforehand to make it easy to work with.

In the process of utilizing and analyzing data, we commonly come across new issues and situations in which a new approach is needed. In such cases, pre-structured data can actually be more difficult to work with, or the structure may be unsuitable to begin with such that the data is not up to task, and at times we have to redo everything from the data structuring step.

Consider a situation in which we find ourselves needing to take our analysis back further in time. The original data will often be archived for storage efficiency reasons, and it will not always be in a state that is easy to work with. Even if it is in a workable state, the degree of processing difficulty rises in proportion to the amount of data.

Preprocessing and structuring the data can reduce the amount of data and increase search speeds, which can be quite beneficial for certain purposes. But what if you could eliminate the time and effort involved in preprocessing, reduce data storage costs, and increase search speeds?

A major advantage of Splunk is that data can be stored in an unstructured, unprocessed form and searched at high speed. Users input raw data without worrying about the data type. Many ways of getting data into Splunk are available—Splunk-specific forwarders are easy to set up, and it can also work with common forwarders like Fluentd.

Even before getting into data utilization requirements, simply putting data into illumino for storage and management already offers decent benefits. It's no exaggeration to say that the act of collecting data itself can create value, because bringing lots of data together and analyzing it as a whole enables the discovery of new value.

© Internet Initiative Japan Inc.

Harnessing high-speed search infrastructure

Data forwarded into illumino (Splunk) is manipulated in various ways to enable high-speed searching. With randomly stored data, search times increase in proportion to data volume, even if ample high-speed storage, CPU, and memory resources are available, and the server costs are also high.

The concept of time series is key here. A lot of data is timestamped. On illumino (Splunk), data is always sorted chronologically.

The statistics show that the searches we need to perform are often on relatively recent data. The most recent data is stored on high-speed SSD storage, and older data is stored on low-cost object storage. illumino (Splunk) returns results seamlessly even when searching data on object storage, so users need not worry about the storage lifecycle. Searches on data in object storage are slightly slower, but the illumino operations team monitors the object storage operating status and tunes the system to achieve optimal data allocation.

We build high-performance search servers with ample CPU and memory in parallel. Depending on the search specifics, a lot of CPU and memory is at times needed to provide the processing power for dealing with unstructured data, but we do maintain sufficient responsiveness for real-world use. Standalone systems (services) are costly, so there are limits to how much server performance is affordable, but common infrastructure allows the costs to be distributed and thus makes abundant infrastructure available. This allows us to use the strengths of Splunk and the cost advantages of common infrastructure to achieve high-speed searching while keeping a lid on the costs borne by each individual user.

Advanced use of Splunk's high-speed searching

For reasonable amounts of data and search details, we are able to achieve sufficient responsiveness using the powerful search capabilities of Splunk products and the illumino environment. But things are less than efficient in some cases when using thousands or tens of thousands of devices within the basic infrastructure to search through huge amounts of data, or when complicated search conditions need to be applied.

In cases like this, you can also structure some of the data imported into Splunk, either when it is imported or at a time of your choosing. The original, raw data is also preserved in these cases, so there is no loss of analytical flexibility. Although the structured data does consume additional storage, the increase in the amount of storage used can be kept relatively small because the system uses something akin to reference pointers to the original data.

Conventional methods often require privileged users—like data administrators—to preprocess data to structure it and so forth, but a big advantage with Splunk is that users can set this up themselves. Clearly, the ability of users to try this out themselves in conjunction with the requirements analysis step will make it possible to achieve the desired output sooner.

Some level of skill with Splunk is needed if users are to take advantage of these Splunk features. At IIJ, therefore,

© Internet Initiative Japan Inc.

support is provided by a dedicated team as a means of reducing learning costs on projects that use the system.

Visualization

Being able to search massive volumes of data at high speed is not the only thing. It is also important that users are presented with results that they can recognize as meaningful data. This is where data visualization comes in.

Extracted data is often presented in a table-like format. Splunk provides "visual effects" that let the user easily transform table data into various types of charts and the like.

In addition to simple line graphs and bar graphs, the user can easily work with dozens of other effects via a GUI. Visualized searches can be executed periodically and presented as a report, and the user can set up a dashboard to display searches on a single screen.

Analysis language

Data search and analysis is accomplished via Splunk's own simple but powerful language, SPL (Search Processing Language). SPL is similar to SQL as used in relational databases in that you specify a data source and filter results according to user-defined conditions, but what makes it powerful is the analysis features that follow that. SPL is a "one-liner" programming language whereby commands are chained using the "|" (pipe) character. Data extracted by a search can simply be passed into an analysis command using a pipe.

For example, say you want to count the number of occurrences of something in chronological order. You can simply add " | timechart count" after the search result. If you wanted to display the number of occurrences for each HTTP status chronologically in a web server log, you would use " | timechart count by status". This is intuitive and simple, so anyone can get up and running with it quickly.

3.3 Challenges and Solutions

The discussion so far has been about the illumino project and its internal system, Splunk. This section looks at challenges and solutions in actual illumino use cases.

3.3.1 Storage, Management, and Visualization of Large Amounts of Data

Challenges

IIJ has many systems running internally, and we use a common system for these systems' network and server infrastructure unless there is any special reason not to. This is the so-called "internal IaaS" approach. The systems are made up of thousands of servers and network devices, and until now, we had been collecting the various logs and metrics in a dedicated system.

This infrastructure has a relatively long history, and although we are able to see information from the logs and metrics using simple visualization tools, we had not been making any further use of or performing any additional analysis on the data. The large scale of the infrastructure means that the amount of data is also large, and reducing the costs involved in storage and management has been an issue we need to address.

Improving data storage and management

The first step was to start bringing this data into illumino. We had already been collecting data using tools like Fluentd, and since the data forwarded to illumino does not need to be structured, it can be sent directly via Fluentd, so all that had to be done was to add illumino as a forwarding destination. That is, we were able to improve cost performance just by adding a simple setting.

Data utilization

To look at the data on conventional systems, you had to view it graphed in a predetermined format on a dedicated system webpage or ask the relevant team to extract the data for you.

© Internet Initiative Japan Inc.

With illumino, authorized users can search the data themselves. They can produce graphs equivalent to those of a conventional system with a few lines of SPL, and since the original data is preserved, they can now set the search period, data category, and the like as they see fit.

A whole range of systems use this internal IaaS offering, and the importance of log and metric analysis varies from system to system. Some are more rigid to data changes while others are more tolerant. Among those that are more rigid, there are those for which CPU load is an important indicator, while disk I/O loads are important on others. Simple analysis and visualizations are provided in common report types and dashboards, and for some projects, custom reports and dashboards that focus on the indicators important for each system are created, providing analysis and visualizations suited to the operations in question.

There are also examples of logs from each system that uses the internal IaaS being imported and the analysis and visualizations being linked with the internal IaaS data to create visualizations of the relationships between each system and infrastructure usage. Visualizing system usage and infrastructure loads in an integrated manner like this helps to improve operational precision.

3.3.2 Data Sharing between Systems/Services

■ Challenges

Customers who use IIJ's services and teams that use internal systems often use multiple systems in combination. They also use combinations of components such as servers and networks, and use combinations such as a business application and a security service. Since individual teams run each of the systems, there was no linking of the data being managed. In the event of a system failure or when conducting surveys/analyses of usage, data had to be looked at across the systems being used, and coordinating the relevant teams took time and effort.

■ Data collection and sharing

As a shared system that can handle unstructured data, illumino is able to bring together data from a range of systems. It offers sufficient advantages just in terms of analyzing data from single systems, but the advantages are even greater when a number of systems and services are used in conjunction with each other.

Say you are using a business application and a security service in combination. Suppose an error appears on the business application screen. And suppose the user reports the details of the error: account name, time of error, etc. Based on that report, the system operator will

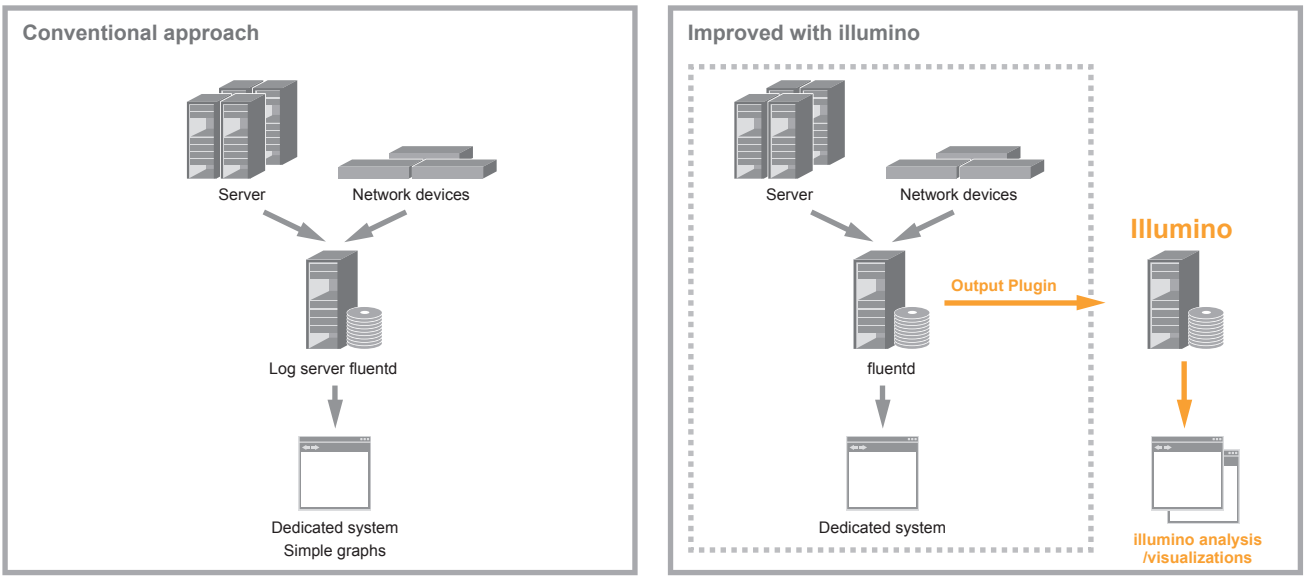

Figure 2: How Data is Utilized with illumino

© Internet Initiative Japan Inc.

start looking through logs and the like to identify the cause. If the business application turns out to be the cause, this is all relatively simple. If you find an error indicating the cause in the logs or whatnot, you can then proceed to the next response step.

What if it's the security service and not the business application that is to blame? Even if the business application log shows an error, this may only go so far as to suggest that the cause lies in the security service, and there may be cases in which no error appears at all. In such cases, the next step would be to investigate the security service side of things. The business processes are designed so that the relevant teams can coordinate smoothly with each other, but some degree of time and effort is still required if this inter-team coordination is not systematized.

Collecting all the data in illumino makes it possible to search and analyze the data in a way that takes into account relationships between datasets. In the above example, the security service can be investigated in conjunction with the investigation of the business application, where the overall investigation started.

Creating an operations tool that links multiple systems in a way that suits the business process design makes it possible to greatly reduce investigative time and effort.

Data security

Although the benefits of data collection and sharing are large, data security also requires careful consideration. Useful data often encompasses sensitive data. It is important to closely examine each system's data security situation and to manage what is forwarded to the shared system and the scope of what is disclosed and shared.

illumino achieves a high level of data security by setting the system up on dedicated internal infrastructure to protect the data and by configuring fine-grained data access permissions.

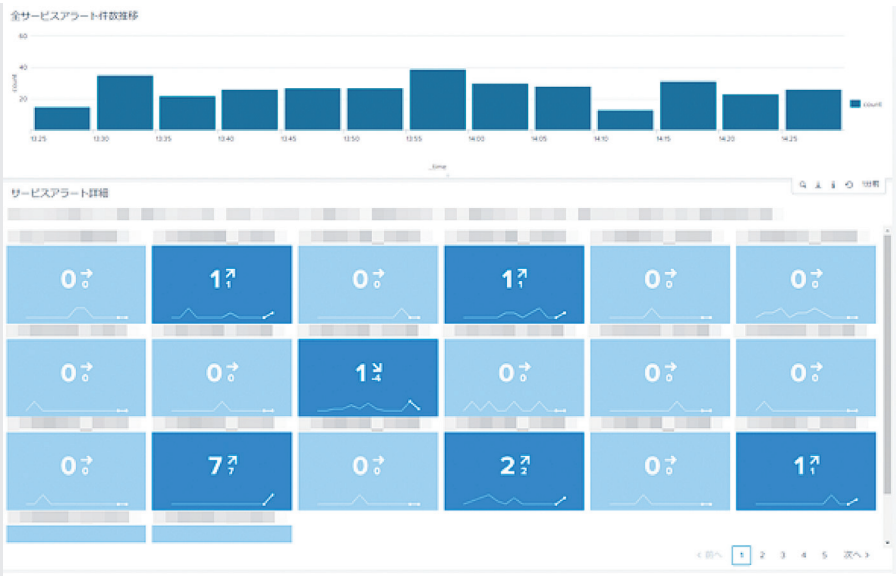

Image of integrated service monitoring dashboard
Multiple services can be monitored on via single screen

Figure 3: Image of Combined Service Monitoring Dashboard

© Internet Initiative Japan Inc.

The permissions need to be managed to a high level of quality, so there is a considerable load and costs involved, but we use role-based permissions management, which facilitates flexible, smart operations in accord with user needs.

3.3.3 Machine Learning

Challenges

One method of converting large amounts of data into useful data is machine learning. It has become relatively easy to acquire machine learning expertise in recent years, and there are real-world examples and solutions on offer, so the need for machine learning within IIJ is also increasing. We found ourselves unable to keep up with this rising need, however, as the task of setting up infrastructure for processing large amounts of data and the time and effort involved in data manipulation do present a high bar.

What is machine learning?

Machine learning is a method of extracting patterns from the original data for prediction and analysis. Analyzing large amounts of complex data using statistically verified methods makes it possible to identify patterns that were difficult to find with a rule-based approach in which operators set hypotheses and performed the design and implementation themselves. The main use cases for the patterns extracted are anomaly detection, prediction, and classification.

One great advantage of Splunk is that it works very well with machine learning by virtue of being a system capable of collecting large amounts of data for high-speed searching and analysis (Figure 4).

Anomaly detection	Predictive analytics	Clustering
	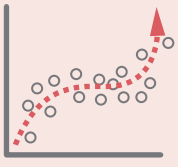	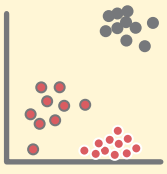
• Deviation from past behavior—Historical AD • Deviation from peers—Cohesive AD • Unusual change in features • Training data available	• Features > predict values—Prediction • Features > predict category—Classification • Time > predict values—Forecasting • Preventive maintenance • Anomaly detection based on deviation from predictions • Training data available	• Identify peer groups—Clustering • Correlation between target variable and feature • Reduce feature granularity • Use rare groups to detect anomalies • Training data unavailable

Figure 4: Anomaly Detection, Prediction, and Clustering via Machine Learning

© Internet Initiative Japan Inc.

Machine learning basics

The machine learning logic itself is not unique to illumino (Splunk). The main machine learning logic provided by Splunk is actually implemented in Python and is open source.

In actuality, you could do machine learning without illumino so long as you have some basic Python skill and knowledge of open source machine learning software. But there are a range of benefits to be had from going through illumino (Splunk).

Advantages of using machine learning with illumino

To perform machine learning, you first need to prepare data to serve as a learning source (learning data). The accuracy of the learning data greatly affects detection quality.

Steps in creating source data suitable for machine learning include:

- Data extraction: Extracting appropriate learning data from a large volumes of original data
- Cleansing: Correcting and filling in noise and missing data
- Data conversion: Normalizing datasets to correspond in terms of data scale

On illumino (Splunk), this preprocessing is done using SPL. The user needs to decide how to extract, correct, fill in, and normalize the data, but once the decisions are made, the processing is easily expressed in SPL. The learning data thus created can be visualized using illumino's graphing features and the like, which is also useful in evaluating accuracy.

Once the learning data is ready, it is fed into the machine learning logic and turned into a model. This process of feeding the data into the machine learning logic is also expressed in SPL. In addition to expressing it in SPL, the user can also feed data into the logic and evaluate accuracy via a GUI. Accuracy can be evaluated visually via the GUI by, for instance, displaying the distribution of the learning data as a graph or using a slide bar to see what changes when different threshold parameters are used. There are data showing that this preprocessing accounts for about 80% of the machine learning process (Figure 5). Using illumino can greatly reduce preprocessing time and effort.

After creating the model, you evaluate it against the test data. The evaluation process can also be expressed in SPL, and the results can be used to issue monitoring alerts, easily visualized in graph and other formats, and

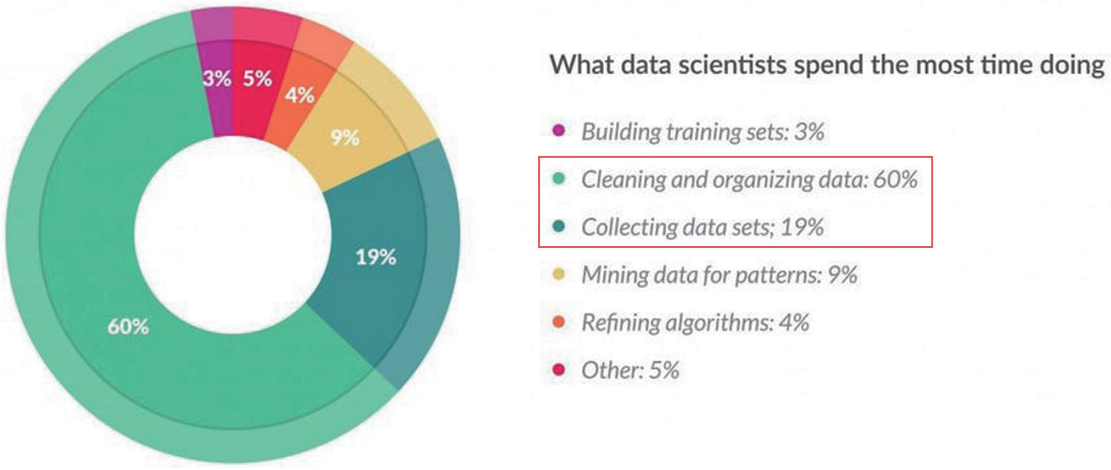

Figure 5: Time Spent on Machine Learning Tasks[*1]

*1 "Cleaning Big Data: Most Time-Consuming, Least Enjoyable Data Science Task, Survey Says", Forbes, March 23, 2016 (https://www.forbes.com/sites/gilpress/2016/03/23/data-preparation-most-time-consuming-least-enjoyable-data-science-task-survey-says/).

© Internet Initiative Japan Inc.

so forth. The ability to complete all steps in the machine learning process on illumino has greatly lowered the bar to its adoption.

Example of machine learning in action

Here, we look at abnormal traffic detection as an example of how IIJ uses the system internally.

On one particular system, we monitor network traffic as one of the system health checks. Previously, the only monitoring configuration we did was to set fixed upper and lower bounds for traffic. As the system is aimed at business users, we know that traffic patterns differ greatly between day and night on weekdays and on holidays. We did at times observe trends in traffic that clearly differed from what was normal, even if not the sort of failure that would cause a complete interruption of service, so there were concerns about some equipment malfunctioning or experiencing abnormalities due to external factors, resulting in certain users being impacted. This was particularly difficult to detect based on fixed thresholds during nighttime and on holidays when traffic levels are relatively low.

To address such cases, we used machine learning to analyze past data and implemented threshold settings based on the probability density distribution (DensityFunction). By performing machine learning on past data on times when the system was operating normally, we can calculate, from a statistical perspective, what we term “rare upper and lower bounds”. By monitoring traffic data with these bounds as the threshold values and conducting a detailed investigation whenever the bounds were exceeded, we were able to identify phenomena that we were previously unaware of.

Fortunately, no failures have occurred since we implemented this machine learning-based monitoring, but if and when failures do occur, we will be able to respond swiftly and with precision.

3.4 Conclusion

We have discussed challenges around the utilization of data at IIJ, the solutions offered by the illumino project, and an example of the system in action and its effects.

The need for data utilization is bound to continue increasing ahead. At the illumino project, we are working to have illumino used on more systems and services and continuously striving to enhance our service offerings.

Takahisa Kudo
Analytics & Management System Development Section, Platform Development Department, Network Division, IIJ
Mr. Kudo joined IIJ in 2008. He is engaged in illumino project planning/operations and serves as illumino system administrator. He is also an internal evangelist for Splunk.

© Internet Initiative Japan Inc.

February 2023 Vol.57

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.

The copyright of this document remains in Internet Initiative Japan Inc. ("IIJ") and the document is protected under the Copyright Law of Japan and treaty provisions. You are prohibited to reproduce, modify, or make the public transmission of or otherwise whole or a part of this document without IIJ's prior written permission. Although the content of this document is paid careful attention to, IIJ does not warrant the accuracy and usefulness of the information in this document.

©Internet Initiative Japan Inc. All rights reserved.
IIJ-MKTG020-0055

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/