IIR

Internet
Infrastructure
Review

Nov. 2025
Vol. 67

Periodic Observation Report (1)

Broadband Traffic Report
Traffic Continues to Grow Steadily

Periodic Observation Report (2)

Effectiveness of Defensive Action and IIJ's Approach to Sender Domain Authentication (ARC)

Focused Research (1)

Answering Calls to Move Away from VMware
—The Tight Relationship Between Kubernetes and Virtual Environments

Focused Research (2)

The 3G Sunset and MVNOs

IIJ
Internet Initiative Japan

Internet Infrastructure Review

November 2025 Vol.67

Executive Summary

The IIR (Internet Infrastructure Review) is a technical journal that IIJ has been publishing quarterly since 2008. From this issue (Vol. 67), Naoshi Someya takes over the writing of the executive summary from Junichi Shimagami. Below is a summary of this issue's contents.

Chapter 1 is our first periodic observation report for this issue. It provides a detailed analysis of trends in traffic on the broadband and mobile services operated by IIJ. From 2024 into 2025, traffic continued to grow steadily overall, with no major changes observed in usage patterns. While the adoption of AI, the use of which continues to grow, has so far had only a limited impact on traffic, this is something that will bear close watching ahead.

Chapter 2 presents our second periodic observation report. It discusses IIJ's defensive action program, which we fully deployed in 2024. This program represents an evolution of traditional abuse countermeasures, enhancing safety for users and society as a whole by proactively detecting and blocking malicious traffic. Defensive action has substantially reduced incidents of spam and phishing attacks while significantly improving the reliability of IIJ's email services. The report also analyzes adoption rates for sender authentication technologies (SPF, DKIM, DMARC, ARC) and transport encryption (STARTTLS), presenting the latest trends both domestically and internationally.

Our first focused research report, in Chapter 3, discusses IIJ's work and challenges faced in relation to reviewing virtualization platforms in the wake of the VMware issue and migrating to next-generation platforms using Kubernetes. In response to rising VMware license prices, IIJ has rolled out an internally developed Kubernetes distribution, IKE (IIJ Kubernetes Engine), as a VMware alternative, successfully reducing costs, improving operational efficiency, and enhancing quality. While Kubernetes was originally a container orchestrator, its use in environments that employ a mixture of containers and VMs according to the use case positions it as a future-ready option capable of generating service value.

Chapter 4 presents our second focused research report, discussing IIJ's initiatives in preparation for the termination of NTT Docomo's 3G services, scheduled for the end of March 2026. The 3G sunset may impact MVNO operators and users, affecting voice calls, SMS, and data communications. IIJ has been identifying technical and operational challenges from an early stage and advancing support and preparations to enable smooth migration to 4G/5G. In particular, the report explains differences in device behavior due to implementation details and provides practical steps for dealing with this.

The IIR aims to provide readers with fresh insights by sharing real voices from the front lines, where IIJ engineers confront challenges every day and help keep the infrastructure that underpins society running. In this era of rapid change, IIJ will continue to take on new challenges and evolve while staying true to its mission of supporting society through technology.

Naoshi Someya

Managing Executive Officer; Network Services Business Unit; Director, Cloud Division, IIJ
Mr. Someya joined IIJ in 1998 and was seconded shortly thereafter to IIJ Technology (which was merged into IIJ in 2010). At IIJ Technology, he was involved in the launch of the systems integration (SI) business and worked on building numerous Internet systems as well as providing consulting services. In 2016, he transferred to IIJ's Service Business Division, where he was responsible for medium-term strategy for the cloud business. In 2019, he became head of the cloud business. As of this fiscal year, he serves as editor-in-chief of the IIR, with his aim being to proactively deliver practical, cross-cutting technical insights from across IIJ to readers.

Broadband Traffic Report
Traffic Continues to Grow Steadily

1.1 Overview

In this report, we analyze traffic over the broadband access services operated by IIJ and present the results each year[*1*2*3*4*5]. Here, we again report on changes in traffic trends over the past year, based on daily user traffic and usage by port.

Overall, traffic continued to grow steadily this year, as it has over the past few years. We see no notable changes in the trends at this point.

Figure 1 plots the overall average monthly traffic trends for IIJ's fixed broadband services and mobile services. IN/OUT indicates the direction from the ISP perspective. IN represents uploads from users, and OUT represents user downloads. Because we cannot disclose specific traffic numbers, we have normalized the data, setting the OUT observations for January 2020, just before the COVID-19 pandemic, for both services to 1.

Over the past year, broadband IN traffic increased 9% and broadband OUT traffic increased 4%. The corresponding year-earlier figures were 14% and 12%, so growth has slowed somewhat. The broadband figures include IPv6/IPoE traffic. IPv6 traffic on IIJ's broadband services comprises both IPoE and PPPoE traffic. As of June 2025, IPoE accounted for a bit under 50% of all traffic, at 42% of IN and 49% of OUT broadband traffic overall. This represents a year-on-year decrease of 1 percentage point for IN and an increase of 1 point for OUT, so the figures are virtually unchanged, suggesting that migration to IPoE has run its course.

Mobile services traffic was largely range-bound in the first year or so of COVID as people went out less, but it has subsequently been in an uptrend. Over the past year, mobile IN traffic increased 29% and mobile OUT traffic increased 9%. The year-earlier figures were 29% and 20%. Mobile services IN traffic accounts for a high proportion of total because of the high volume of uploads on services for enterprise customers. Looking solely at personal services, IN accounts for around a tenth of the total, similar to the situation for broadband.

We now look at broadband traffic by time of day on weekdays over the past year. Figure 2 plots hourly average

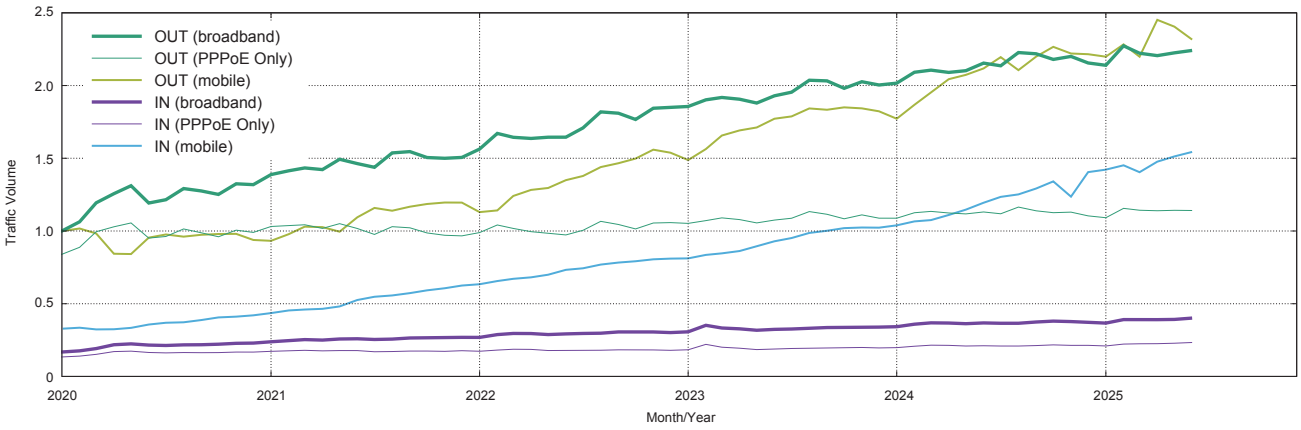

Figure 1: Monthly Broadband and Mobile Traffic

*1 Kenjiro Cho. Broadband Traffic Report: Looking Back on the Past 5 Years. Internet Infrastructure Review. Vol. 64. pp4–11. September 2024.
*2 Kenjiro Cho. Broadband Traffic Report: Traffic in a Stable Uptrend Post-COVID. Internet Infrastructure Review. Vol. 60. pp4–9. September 2023.
*3 Kenjiro Cho. Broadband Traffic Report: COVID's 3rd Year Brings Lull in Traffic. Internet Infrastructure Review. Vol. 56. pp4–11. September 2022.
*4 Kenjiro Cho. Broadband Traffic Report: COVID-19's Impact in its 2nd Year. Internet Infrastructure Review. Vol. 52. pp4–11. September 2021.
*5 Kenjiro Cho. Broadband Traffic Report: The Impact of COVID-19. Internet Infrastructure Review. Vol. 48. pp4–9. September 2020.

traffic volume for Monday–Friday for four one-week blocks selected at intervals of roughly four months since early June 2024. Weekday daytime traffic volumes have increased during school holiday periods in recent years, so we selected school weeks. Traffic volume here is the sum of PPPoE and IPoE. The dotted lines in the lower part of the plot represent uploads for each week, but focusing again on download volumes in this edition, we see that while traffic volumes have not increased much for the hours from late night through early morning, they have steadily increased during the hours from morning through evening.

1.2 About the Data

As with previous reports, for broadband traffic, our analysis uses data sampled using Sampled NetFlow from the routers that accommodate the fiber-optic and DSL broadband customers of our personal and enterprise broadband access services. For mobile traffic, we use access gateway billing information to determine usage volumes for personal and enterprise mobile services, and we use Sampled NetFlow data from the routers used to accommodate these services to determine the ports used.

Because traffic trends differ between weekdays and weekends, we analyze traffic in one-week chunks. In this report, we look at data for the week of June 2–8, 2025, and compare those data with data for the week of June 3–9, 2024, which we analyzed in the previous edition of this report.

Results are aggregated by subscription for broadband traffic, and by phone number for mobile traffic as some subscriptions cover multiple phone numbers. The usage volume for each broadband user was obtained by matching the IP addresses assigned to users with the IP addresses observed. Note that IPoE traffic is not included in the analysis of traffic by port, as detailed data are not available because we use Internet Multifeed Co.'s transix service for IPoE.

1.3 Users' Daily Usage

First, we examine daily usage volumes for broadband and mobile users from several angles. Daily usage indicates the average daily usage calculated from a week's worth of data for each user.

Since our 2019 report, we have used daily usage data only on services provided to individuals. The distribution is heavily distorted if we include enterprise services, where usage patterns are highly varied. So to form a picture of overall usage trends, we determined that using only the personal user data would yield more generally applicable, easily interpretable conclusions. Note that the analysis of usage by port in the next section does include enterprise data because of the difficulty of distinguishing between individual and enterprise usage. Note also that we have included IPoE user data in the broadband figures since 2021, so the broadband data comprise both PPPoE and IPoE*6.

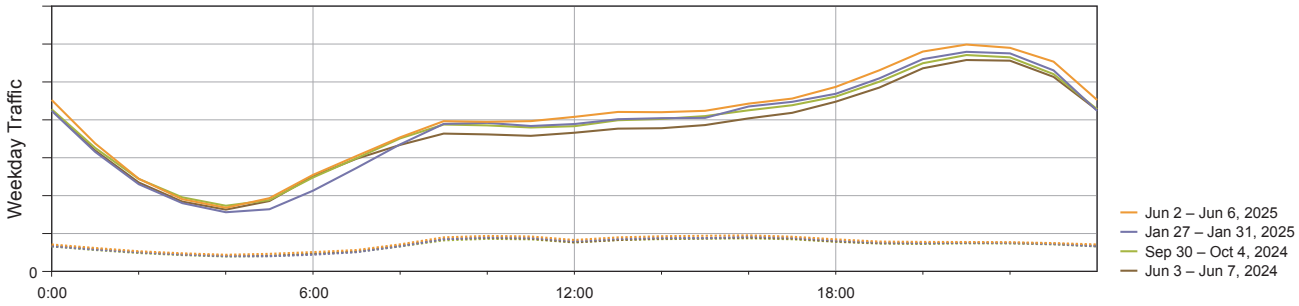

Figure 2: Hourly Average Broadband Traffic on Weekdays in the Past Year

*6 The PPPoE and IPoE usage figures of users who use both protocols are treated as coming from separate users.

Figures 3 and 4 show the average daily usage distributions (probability density functions) for broadband and mobile users. Each compares data for 2024 and 2025 split into IN (upload) and OUT (download), with user traffic volume plotted along the X-axis and user frequency along the Y-axis. The X-axis shows volumes between 10KB (104) and 1TB (1012) using a logarithmic scale. Most users fall within the 1TB (1012) range, with a few exceptions.

The IN and OUT traffic distributions in the figures are close to a log-normal distribution, which looks like a normal distribution on a semi-log plot. A linear plot would show a long-tailed distribution, with the peak close to the left. The OUT distribution is further to the right than the IN distribution, indicating that download volume is more than an order of magnitude larger than upload volume.

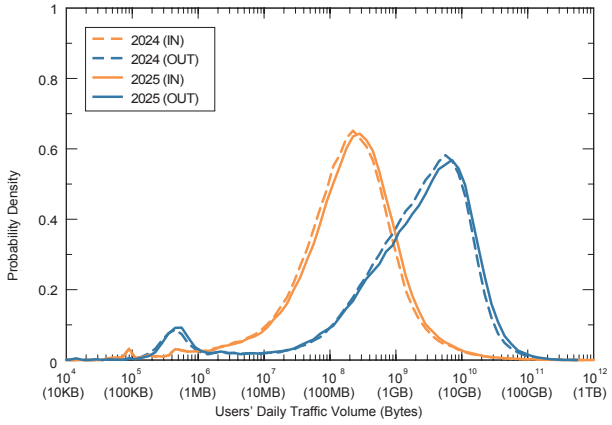

Figure 3: Daily Broadband User Traffic Volume Distribution Comparison of 2024 and 2025

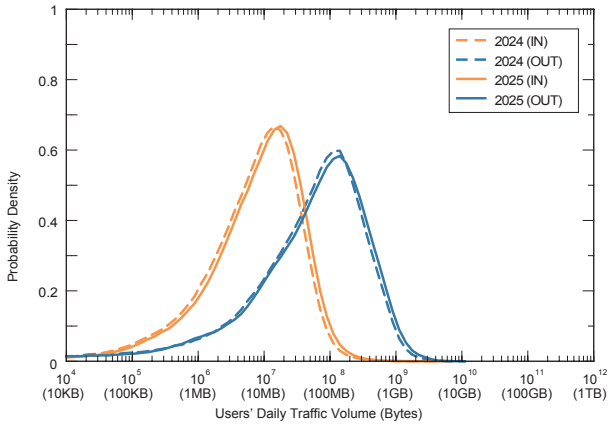

Figure 4: Daily Mobile User Traffic Volume Distribution Comparison of 2024 and 2025

First, we look at the broadband distributions in Figure 3. Comparing 2024 and 2025, both the IN and OUT distributions have moved slightly to the right, indicating that overall traffic volume has increased.

The peaks of the mobile distributions in Figure 4 have also moved a little to the right since last year, indicating that overall traffic has increased. Mobile usage volumes are significantly lower than for broadband, and limits on mobile data usage mean that heavy users, which fall on the right-hand side of the distribution, account for only a small proportion of the total. There are also no extremely heavy users. The variability in each user's daily usage volume is higher for mobile than for broadband owing to there being users who only use mobile data when out of the home/office as well as limits on mobile data.

Table 1 shows trends in the mean and median daily traffic values for broadband users as well as the mode (the most frequent value, which represents the peak of the distribution). When the peak is slightly off the center of the distribution, the mode is adjusted to bring it

	IN (MB/day)			OUT (MB/day)		
Year	Mean	Median	Mode	Mean	Median	Mode
2007	436	5	5	718	59	56
2008	490	6	6	807	75	79
2009	561	6	6	973	91	100
2010	442	7	7	878	111	126
2011	398	9	9	931	144	200
2012	364	11	13	945	176	251
2013	320	13	16	928	208	355
2014	348	21	28	1124	311	501
2015	351	32	45	1399	443	708
2016	361	48	63	1808	726	1000
2017	391	63	79	2285	900	1259
2018	428	66	79	2664	1083	1585
2019	479	75	89	2986	1187	1995
2020	609	122	158	3810	1638	3162
2021	714	143	200	4432	2004	3981
2022	727	142	178	4610	2010	3981
2023	804	166	224	5456	2369	5012
2024	834	178	224	5743	2372	5620
2025	886	202	282	6538	2615	6310

Table 1: Trends in Mean and Mode of Broadband Users' Daily Traffic Volume

toward the center. Comparing 2024 and 2025, the IN mode rose from 224MB to 282MB while the OUT mode rose from 5,620MB to 6,310MB, translating into growth factors of 1.26 for IN and 1.12 for OUT. Meanwhile, because the means are influenced by heavy users (on the right-hand side of the distribution), they are significantly higher than the corresponding modes, with the IN mean at 886MB and the OUT mean at 6,538MB in 2025. The 2024 means were 834MB and 5,743MB, respectively. As mentioned, up to 2020 the data covered only PPPoE users, and since 2021 the data have covered both PPPoE and IPoE users.

Table 2 shows the mobile traffic metrics. In 2025, the IN mode was 16MB and the OUT mode was 126MB, while the means were IN 19MB and OUT 172MB. The 2024 modes were IN 14MB and OUT 112MB, and the means were IN 16MB and OUT 150MB.

	IN (MB/day)			OUT (MB/day)		
Year	Mean	Median	Mode	Mean	Median	Mode
2015	6.2	3.2	4.5	49.2	23.5	44.7
2016	7.6	4.1	7.1	66.5	32.7	63.1
2017	9.3	4.9	7.9	79.9	41.2	79.4
2018	10.5	5.4	8.9	83.8	44.3	79.4
2019	11.2	5.9	8.9	84.9	46.4	79.4
2020	10.4	4.5	7.1	79.4	35.1	63.1
2021	9.9	4.7	7.9	85.9	37.9	70.8
2022	12.8	6.0	10.0	113.7	49.2	89.1
2023	14.1	6.8	11.2	129.2	56.0	100.0
2024	16.3	8.2	14.1	150.4	66.7	112.2
2025	19.3	9.7	15.8	172.2	73.5	125.9

Table 2: Trends in Mean and Mode of Mobile Users' Daily Traffic Volume

Figures 5 and 6 plot per-user IN/OUT usage volumes for random samples of 5,000 users. The X-axis shows OUT (download volume) and the Y-axis shows IN (upload volume), with both using a logarithmic scale. Users with identical IN/OUT values fall on the diagonal.

The cluster spread out below and parallel to the diagonal in each of these plots represents typical users with download volumes an order of magnitude higher than upload volumes. Variability between users in terms of usage levels and IN/OUT ratios is wide, indicating that there is a diverse range of usage styles. For mobile traffic, the pattern of OUT being an order of magnitude larger also applies, but usage volumes are much lower than for broadband. For both broadband and mobile, there appears to be almost no difference between these plots and those for 2024.

Traffic is heavily skewed across users, such that a small proportion of users accounts for the majority of overall traffic volume. For example, the top 10% of broadband users account for 50% of total OUT and 73% of total IN traffic, while the top 1% of users account for 15% of OUT and 44% of IN traffic. On mobile, the top 10% of users account for 48% of total OUT and 46% of total IN traffic, while the top 1% of users account for 12% of OUT and 13% of IN traffic. These proportions have hardly changed from last year.

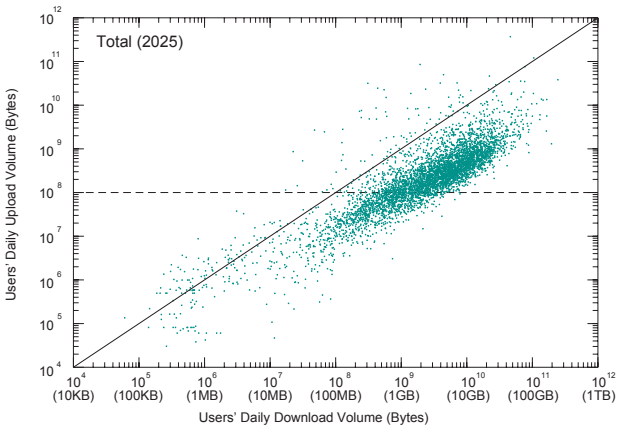

Figure 5: IN/OUT Usage for Each Broadband User

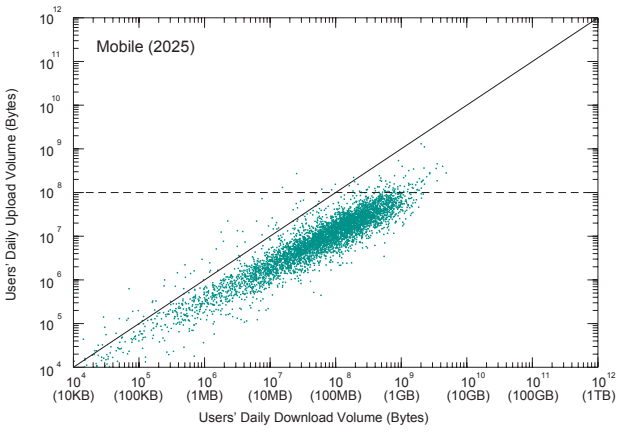

Figure 6: IN/OUT Usage for Each Mobile User

1.4 Usage by Port

Next, we look at a breakdown of traffic and examine usage levels by port. Recently, it has become difficult to identify applications by port number. Many P2P applications use dynamic ports on both ends, and a large number of client/server applications use HTTP ports like port 80 to avoid firewalls. Hence, generally speaking, when both parties are using a dynamic port numbered 1024 or higher, the traffic is likely to be from a P2P application, and when one of the parties is using a well-known port lower than 1024, the traffic is likely to be from a client/server application. In light of this, we take the lower of the source and destination port numbers when breaking down TCP and UDP usage volumes by port.

Table 3 shows the percentage breakdown of broadband users' usage by port over the past five years. In 2025, 66% of all traffic was over TCP connections, down 2 points from 2024. The proportion of traffic over port 443 (HTTPS) was 53%, only a slight drop from 2024. The proportion of traffic over port 80 (HTTP) was 6%, down 1 point. The figure for UDP port 443, which is used by the QUIC protocol, was up 2 points to 23%.

TCP dynamic port traffic fell ever so slightly to 6%. Individual dynamic port numbers account for only a tiny portion, with the most commonly used port 31000 only making up 1.6%.

Table 4 shows the percentage breakdown by port for mobile users. The figures are close to those for broadband on the whole. This is possibly because apps similar to those for PC platforms are now also used on smartphones, and because the proportion of broadband usage on smartphones is rising.

The broadband port data only include PPPoE, not IPoE, and so do not necessarily reflect the trend in fixed broadband overall. Comparing IPv4 and IPv6 on mobile, port 443 accounts for a higher proportion of both TCP

year	2021	2022	2023	2024	2025
protocol port	(%)	(%)	(%)	(%)	(%)
TCP	**71.9**	**71.6**	**70.5**	**67.5**	**65.5**
(< 1024)	65.8	65.4	64.8	61.1	59.8
443 (https)	53.5	55.7	56.9	53.8	53.2
80 (http)	11.6	8.9	7.2	6.5	5.9
993 (imaps)	0.1	0.1	0.1	0.1	0.2
183	0.1	0.2	0.2	0.2	0.1
22 (ssh)	0.2	0.1	0.1	0.1	0.1
(>= 1024)	6.1	6.2	5.7	6.4	5.7
31000	0.6	0.9	1.1	1.2	1.6
8080	0.4	0.3	0.4	0.3	0.3
1935 (rtmp)	0.2	0.2	0.2	0.3	0.2
UDP	**24.5**	**24.3**	**25.4**	**28.2**	**30.6**
443 (https)	15.9	16.3	18.2	21.0	23.1
4500 (nat-t)	0.8	0.8	1.0	0.9	0.7
8801	0.9	0.6	0.4	0.4	0.3
ESP	**3.3**	**3.8**	**3.8**	**4.0**	**3.6**
GRE	**0.2**	**0.2**	**0.1**	**0.2**	**0.2**
IP-ENCAP	**0.1**	**0.1**	**0.1**	**0.1**	**0.1**
ICMP	**0.0**	**0.0**	**0.0**	**0.0**	**0.0**

Table 3: Broadband Users' Usage by Port

year	2021	2022	2023	2024	2025
protocol port	(%)	(%)	(%)	(%)	(%)
TCP	**70.3**	**71.6**	**71.0**	**71.0**	**69.8**
443 (https)	44.4	42.3	42.1	42.2	37.8
80 (http)	5.0	4.1	3.5	1.8	1.5
993 (imaps)	0.2	0.1	0.1	0.1	0.1
1935 (rtmp)	0.1	0.1	0.2	0.1	0.1
UDP	**23.8**	**24.4**	**26.5**	**27.5**	**29.3**
443 (https)	16.3	17.9	20.9	22.5	24.8
4500 (nat-t)	3.7	2.7	2.5	1.8	1.5
51820	0.0	0.1	0.2	0.3	0.3
53 (dns)	0.2	0.2	0.2	0.2	0.2
8801	0.7	0.3	0.2	0.1	0.1
ESP	**5.8**	**3.9**	**2.4**	**1.4**	**0.8**
ICMP	**0.0**	**0.0**	**0.1**	**0.0**	**0.1**
GRE	**0.1**	**0.0**	**0.0**	**0.0**	**0.0**

Table 4: Mobile Users' Usage by Port

and UDP usage on IPv6, and there is probably a similar trend in the case of IPoE.

Figure 7 compares overall broadband traffic for key port categories across the course of the week from which observations were drawn in 2024 and 2025. We break the data into four port buckets: TCP ports 80 and 443, dynamic TCP ports (1024 and up), and UDP port 443. The data are normalized so that peak overall traffic volume on the plot is 1. The overall peak is around 19:00–23:00. There are no major changes overall relative to 2024, but traffic on UDP port 443 increased a little.

Figure 8 shows the trend for TCP ports 80 and 443 and UDP port 443, which account for the bulk of mobile traffic. As was the case with broadband, mobile traffic on UDP port 443 was up slightly compared with 2024. Comparing the plots with those for broadband, usage times evidently differ, with mobile having three separate traffic peaks on weekdays: morning commute, lunch break, and evening.

We now examine weekday time-of-day traffic volumes for major content providers. Figure 9 plots broadband TCP and UDP traffic with source port 443. We mapped source IP addresses to their AS numbers (ASNs), using this to identify the registrant organization, and plotted weekday time-of-day traffic for major providers' ASNs at two-hour intervals. Many major providers operate multiple ASNs, but for each provider we include only the ASN with the largest traffic volume. These providers also deliver content via third-party content delivery networks (CDNs),

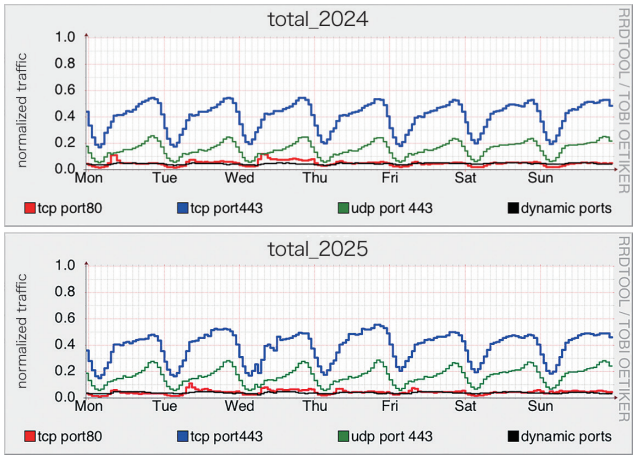

Figure 7: Broadband Users' Port Usage Over a Week
2024 (top) and 2025 (bottom)

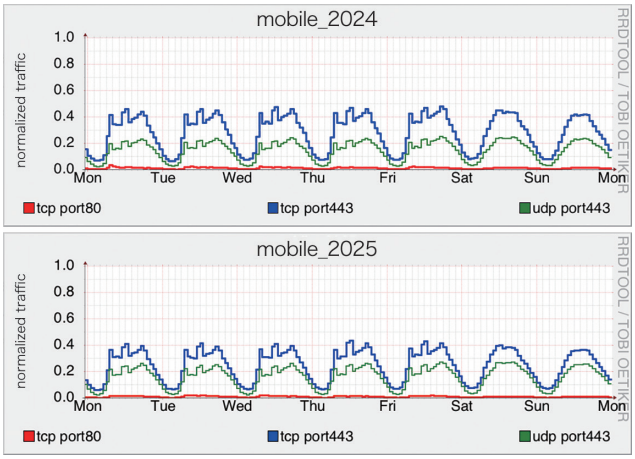

Figure 8: Mobile Users' Port Usage Over a Week
2024 (top) and 2025 (bottom)

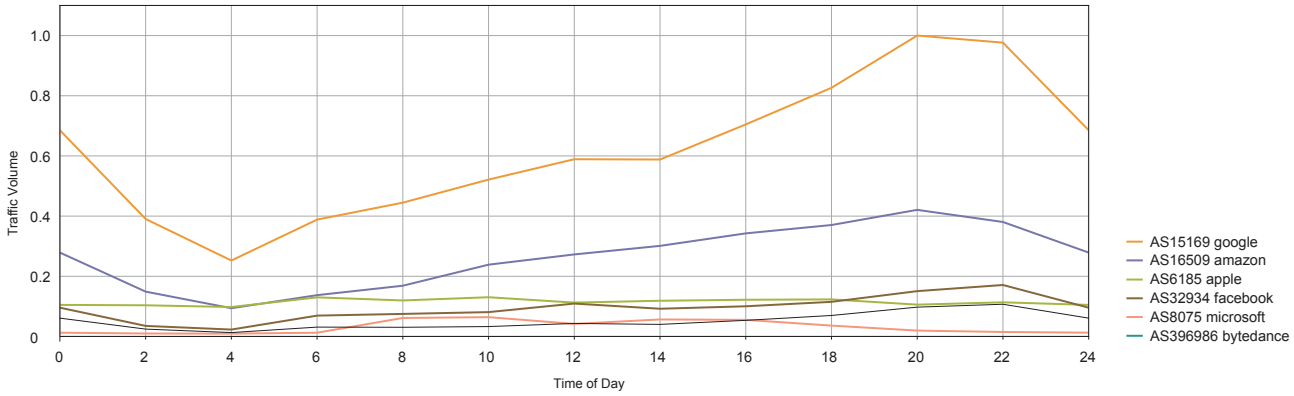

Figure 9: Major Content Providers' Hourly Average Broadband Traffic on Weekdays

and as we are unable to capture and count traffic served through such CDNs the volumes shown do not represent each provider's total traffic; they are strictly per-ASN figures.

Google, which includes YouTube, has the highest traffic volume. Second is Amazon, which includes Amazon Prime Video as well as traffic from other companies using Amazon Web Services. Google, Amazon, Facebook, and ByteDance, which operates TikTok, are all video-heavy providers, with traffic peaking from evening into night, reflecting household video viewing patterns. By contrast, Apple's traffic remains steady throughout the day, likely due to automatic app updates and the like. Microsoft sees more traffic during daytime hours, suggesting work-related usage such as for remote work. Looking only at UDP port 443, Google accounts for 67%, Facebook for 13%, and ByteDance for 6%.

1.5 Conclusion

Broadband traffic has been growing relatively steadily over the past few years, with very little change in the overall trend. But looking back, we have repeatedly seen several years of little change in traffic then being followed by the next shift, so the next wave might be on the way soon.

One possible such wave could come from the currently hot topic of AI, but even if AI usage among users in general rises, the impact on broadband traffic is likely to be limited. We can expect, for instance, users to shift from keyword searches to AI chatbots, but this would not entail any major change in data size, so there would not be much impact on traffic volume.

Another possibility is a scenario in which a major terrorist incident or similar event sparks the rapid adoption of

cloud-based AI services for analyzing surveillance camera footage, resulting in a sharp rise in upload traffic. While such services are likely to expand ahead, rapid growth may be unlikely due to privacy concerns about providing surveillance camera footage to third parties.

Behind the scenes, some organizations developing large-scale AI models appear to be fetching vast amounts of content for use as AI training data, which in turn is increasing load and causing other impacts for content providers.

Broadband traffic growth is currently driven by increases in data-heavy video content. In recent years, watching Internet video has become smooth and hassle-free both at home and on mobile; you can see many people watching videos on their smartphones even on the train. Even if the pace of traffic growth does not change markedly, the number of video users, per-person viewing time, and the data volume associated with higher video quality are all likely to continue increasing for the time being.

These are merely quantitative changes, but the fact that anyone can now shoot video on a smartphone and easily edit and share it is a qualitative change, and as AI makes working with video even easier, rapid progress can be expected. This phenomenon goes beyond a mere increase in online video viewing. We see it as a broader social phenomenon whereby technology is fundamentally changing the way we communicate, and indeed our culture itself, with the shift from a text-centric culture to one of photos plus short posts on social media, and onward to video and animation.

Kenjiro Cho
Research Director, Research Laboratory, IIJ

Effectiveness of Defensive Action and IIJ's Approach to Sender Domain Authentication (ARC)

2.1 New Initiatives to Protect Our Customers: Effectiveness of Defensive Action

2.1.1 Building on the Previous Report

In our previous report, we discussed the issue of malicious actors hijacking email accounts and exploiting email services to send, for example, phishing emails.

When email services are misused by malicious actors, not only are phishing email recipients targeted by attacks but other customers also using the email service are impacted by reduced availability of service infrastructure, reduced deliverability to destination email services, and the like. This sort of unwelcome and fraudulent behavior on the Internet occurs not only on IIJ's services but also on those of other ISPs and third parties on a daily basis. It is commonly referred to as abuse, and a challenge has been that only post-incident responses were possible.

As such, we launched a new initiative on May 1, 2024, under the IIJ Secure MX Service contract terms in an effort to maintain email service quality and protect our customers[*1]. This initiative involves detecting preparations for improper use and restricting communications to the extent necessary before phishing emails are actually sent out to prevent abuse from occurring in the first place.

2.1.2 Effectiveness of Defensive Action

While abuse responses involve taking action against abuse that has already occurred, we refer to this new initiative as defensive action because it protects customers by detecting preparations for abuse ahead of time.

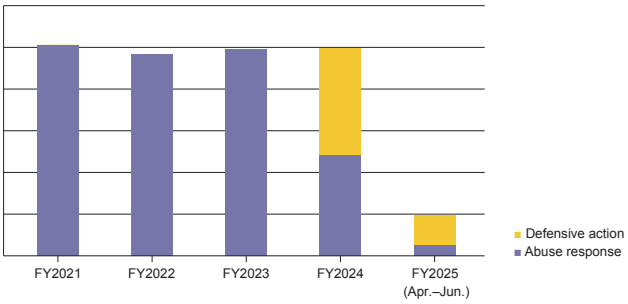

Figure 1: Comparison of abuse response and defensive action incident counts

Roughly a year has now passed since we initiated this defensive action program, and in this article, we report on the effects it has had.

Figure 1 is a stacked bar chart aggregating the number of abuse response and defensive action cases. The vertical axis represents the number of abuse or defensive incidents, while the horizontal axis represents fiscal years (April of one year to March of the following year). For reference, we also include data going back three years before (FY2021) we commenced defensive action.

As is evident from the graph, while the number of abuse response incidents was virtually constant over the previous three years, our defensive action starting in FY2024 has roughly halved the number of abuse response cases.

As explained at the beginning, when abuse responses occur, email service quality is compromised and engineers must undertake unplanned work. Through defensive action, we have been able to eliminate around 50% of these factors.

The fact that defensive action can restrict the transmission of emails in advance and thus prevent abuse means that we are suppressing the transmission of phishing emails and the like from IIJ's network. With email, the sender and recipient perspectives are two sides of the same coin, so defensive action can be regarded as a technical initiative that supports the stable operation of Internet infrastructure and contributes to improved reliability, thus making it highly effective.

While the FY2025 data only go up to June due to publication timing, even more interesting findings have emerged. While we saw roughly equal proportions of abuse response and defensive action during FY2024, the three months of data for FY2025 indicate that the proportion of defensive action cases has increased relative to abuse response cases.

*1 For a detailed background and information on this initiative, see IIR Vol.63 (https://www.iij.ad.jp/dev/report/iir/063/01.html).

Making IIJ's email service infrastructure difficult for malicious actors to exploit can be expected to gradually deter abuse. From a service product owner's perspective, as the number of abuse response cases decreases, engineer resources can be redirected toward other initiatives such as operational improvements, quality enhancements, and customer support.

At IIJ, we will continue working to protect our customers going forward.

2.2 IIJ's Approach to Sender Authentication (ARC)

2.2.1 Background and Overview of Our Approach

Following Google's and Yahoo's 2023 announcements that they would be requiring sender authentication (SPF, DKIM, DMARC), IIJ took steps to ensure its internal systems were compliant.

As part of this process, we discovered that ARC (Authenticated Received Chain) authentication was failing when forwarding to Microsoft 365 (M365). ARC plays a role in preventing false positives under DMARC policies by reevaluating SPF and DKIM authentication results during email forwarding and adding those results and signature information to the headers. The email headers at time of verification contained the string arc=fail (47), as shown below, indicating an authentication failure.

```
ARC-Authentication-Results: i=2; mx.microsoft.com 1; spf=fail (sender ip is
...) smtp.rcpttodomain=iijsmxstg.onmicrosoft.com
smtp.mailfrom=iij.ad.jp; dmarc=pass (p=reject sp=reject pct=100) action=none
header.from=iij.ad.jp; dkim=pass (signature was verified) header.d=iij.ad.jp;
arc=fail (47)
```

To determine whether this issue originated from IIJ's systems or from processing performed on other email systems, we conducted ARC verification tests with multiple providers.

Table 1 summarizes the forwarding test results for the providers involved. IIJ-office is the email system used internally at IIJ for business operations, while SMX is the IIJ Secure MX Service, a different email system.

2.2.2 Inquiring About the Issue

To identify the cause of the ARC authentication failures, we performed signature verification using dkimpy, a Python library that complies with RFC 6376 and RFC 8617. This confirmed that ARC-Message-Signature (AMS) verification was succeeding for emails from IIJ to M365.

```
>>> import dkim
>>> dkim.ARC(open("iijsmx-forward-365-arc-fail-20231127.eml","br").read()).verify()
```

```
(b'fail', [{'instance': ..., ...; spf=... smtp.rcpttodomain=... smtp.mailfrom=...;
dmarc=... action=... header.from=...; dkim=... header.d=...; arc=fail (47)\r\n',
'ams-domain': ..., 'ams-selector': ..., 'ams-valid': True, 'as-domain': ...,
'as-selector': ..., 'cv': ..., 'as-valid': ...}], "x= ...")
```

Given these results and our growing suspicions, we contacted Microsoft. The response was that arc=fail was occurring because hash values did not match during ARC signature verification, suggesting the possibility that the message had been modified after signing.

Before the hash value is calculated, the email content is normalized. Our investigation revealed that while IIJ was using the simple algorithm, M365 was using the relaxed algorithm.

Normalization is the process of converting email content into a standardized format, in compliance with RFC 6376, the DKIM standard. RFC 6376 defines two normalization algorithms for headers and body content: the "simple" and the "relaxed" algorithm. The simple algorithm keeps line breaks and other whitespace as is, faithfully preserving the content as sent. In contrast, the relaxed algorithm ignores line breaks and other whitespace and consolidates all whitespace sequences into a single space character.

RFC 8617, the ARC standard, explicitly states that ARC follows DKIM syntax and processing. For the bh (body hash) tag, in particular, it requires hash values to be calculated on the normalized body content, just as with DKIM.

Forwarding path	i=1	i=2	i=3	ARC result
IIJ-office → SMX → Google → FastMail	IIJ-office	SMX	Google	fail
IIJ-office → SMX → Google → Microsoft	IIJ-office	SMX	Google	fail
IIJ-office → SMX → FastMail	IIJ-office	SMX	-	pass
IIJ-office → SMX → Microsoft	IIJ-office	SMX	-	fail

Table 1: Forwarding Test Results for Each Provider

We also informed Microsoft that the hash values generated by IIJ services matched those from Gmail and OSS implementations, and that M365's normalization algorithm might differ from what other services and OSS use. But because we could not completely rule out the possibility that IIJ's email systems were the cause of the arc＝fail, Microsoft said that investigating the issue would be difficult and thus closed the inquiry.

2.2.3 Reinvestigating and Identifying the Issue

Subsequently, in response to arc＝fail incidents identified in other cases, we corrected issues on the IIJ side and, as a temporary measure, switched from simple to relaxed for the body normalization algorithm. We had expected this change to resolve the ARC authentication failure issue, but we continued to observe incidences of arc＝fail, prompting another detailed investigation.

As with our first inquiry, we conducted tests of forwarding to other providers while examining the values in the email headers more closely. When we tested both empty-body emails and emails containing text as specimens, we found that for empty-body emails, the bh tag hash values differed between IIJ and M365. IIJ requested that Microsoft correct this issue.

```
ARC-Message-Signature: i=3; a=rsa-sha256; c=relaxed/relaxed; d=microsoft.com;
 s=arcselector9901;
 h=From:Date:Subject:Message-ID:Content-Type:MIME-Version:X-MS-Exchange-
AntiSpam-MessageData-ChunkCount:X-MS-Exchange-AntiSpam-MessageData-0:X-MS-
Exchange-AntiSpam-MessageData-1;
 bh=47DEQpj8HBSa+/TImW+5JCeuQeRkm5NMpJWZG3hSuFU=;

b=P0BjQ+kWDc+ghsJ3atV7ElnsRM8v6qSXVbWk+sT59USGlCYId0vvOL5H0s5sfLj58Ow4HnWmYG9RN8As9
rDEBb3IlWihg13+kPDEM8/S3HydFI58bT3KjHuWzM5UGH3MBv+lriGKciDpbDRhnlaz2mvZNnBMuBqCD0eD
8NAa3soY6oP8/jh6t692dWiDRD9pI+D8ho9VhpRZw1SFQ7UIJAtWMLNuY6YG9WlGXmdSi9nCALdhBzwGjod
1Xx0/+RnxLZQAg56eSYYAqCTC5At/YbCjK8b6Yk7+MLcuJBAL/JYAy69uJiBcNo3sroAbnjp7zhgi4FvILn
oM3pMvzFidnA==

ARC-Message-Signature: i=2;a=rsa-sha256;c=relaxed/relaxed;d=securemx.jp;h=
      Content-Type:MIME-Version:Subject:Message-Id:To:From:Date:DKIM-Signature;s=
      arc20190303;t=1713333099;x=1713937899;bh=frcCV1k9oG9oKj3dpUqdJg1PxRT2RSN/XK
      dLCPjaYaY=;b=XT741y7ourI4cRnEqfSt2iNOXIsKAbHFNGzRrMz0R/Rfnm/UuewENteicKQbQZ
      A5E6dfKj1osLQvWGSIwUDQjb405Y0yt3PS/rrV4Eb+LmfvXLGdbVCUT7Q8HPchkyIXgy6D2LEB3
      eF69Br0MF97GW2jz5YF9Dj51I6VuytabrDa31B1w62ENmz7N/gTT62aAPYLUBuHL9DhuJi1d9lm
      +uZLh3ia8iswz8FxZDkuRbrBwQu8eyzinlPyqAfUdPcTbERf+GuGDOfkFyL2NHUg4HUHGaaoqQI
      yt3Vuqt31eM/Uz75LGcPITpRWPGUxZviMLTMiV/Jeoxd4CA2+wbEcOg==

ARC-Message-Signature: i=1;a=rsa-sha256;c=relaxed/relaxed;d=securemx.jp;h=
      Content-Type:MIME-Version:Subject:Message-Id:To:From:Date:DKIM-Signature;s=
      arc20190303;t=1713333098;x=1713937898;bh=frcCV1k9oG9oKj3dpUqdJg1PxRT2RSN/XK
      dLCPjaYaY=;b=T0diE3VZQzYgPWNPb3mPL/hc9Bsw/xc2LOvdgzq5D/P19MlMswEm7oO7a1pKgr
      EvtHRelRPQFCPFP4p5Fc6/CFITiMjRBBDLaqBguN3VsbjwyRbp1BSTzFzEm8+/2hI6hFri6XwfY
      TmfiE799hedDK5XoFCjURuT/gv4MSFjdTzLkdZY2M4CY00fHPiM9g7av9aJU3OYfku8DUrFev4t
      dIWI/7jT6tRpsN/roh3WEb1j3ll9YUqTMMOhN4fzO2Teo+4BCLxjGVZLL60iglOA8tkO899QtCo
      aDJqM8BeSWFc0m8WO1CYTg92HukBJQphxb3g8XMayEB0n6B3sQ20E0A==
```

When text was present, the bh tag hash values appeared to be the same.

```
ARC-Message-Signature: i=3; a=rsa-sha256; c=relaxed/relaxed; d=microsoft.com;
 s=arcselector9901;
 h=From:Date:Subject:Message-ID:Content-Type:MIME-Version:X-MS-Exchange-
AntiSpam-MessageData-ChunkCount:X-MS-Exchange-AntiSpam-MessageData-0:X-MS-
Exchange-AntiSpam-MessageData-1;
 bh=KXMY6AzIdALlOzgkBat5CPcMk0Vqq3IozVvMgC0v4D0=;

b=hoDE5ZJ2ZnJuUTrPAsRdaZjjfsIHGRAH16O7Syq0uNPc7MPkKWyCUaRhU5JR14hg+TIClmfdauNL9nLN6
MDPrdAWBGwQr00h8fQrcy4AHrK206Cex9YV07/AjBu7eO09ld7wTWr3IwJc8wzQ83CYXQ4AIoXqK+mzsXo8
rQDOaPOxgTmNwvfGm+Q0Q5AdNzoesniZdAfmC1ZjKw774bezqKlTEY9P9ylwQPyQ830nHpbToug4PlQN5la
QReb2wxFVQUzfY8Wser01/31hzuM1kB1Regz9o9+gRBJiMs3YBP4ENLUOqp0Oq7dUhIDhZ9T++hGHvNZtZN
Ckheo4HIClOw==

ARC-Message-Signature: i=2;a=rsa-sha256;c=relaxed/relaxed;d=securemx.jp;h=
      Content-Type:MIME-Version:Subject:Message-Id:To:From:Date:DKIM-Signature;s=
      arc20190303;t=1713333207;x=1713938007;bh=KXMY6AzIdALlOzgkBat5CPcMk0Vqq3IozV
      vMgC0v4D0=;b=zLi0vrLrkhYhpkGqnKJ2IyYywsXPSc3vGsVUjV5UOzuX+Pg3VdgwdXgXDbANw9
      PdpDJ/uBklbyU0v52Vw7Sa8RPiB3QdaZOfLoIyEa3uDAGqy7cKd+yAXfnkBGQoZhP8ddkT8yQZz
      uK9/Vz0F5tQb+RTzkjjugHuavng6OKA1YTCIYPB654uqicxyx+L6TwLzajKv20TKZIDYq6fpQFO
      4vW0bclHa4hKk91lkI7yAPZ6WpuhZv6S9ySnIs6PujDSXYV2GEMcu8aTa4VHAXNGLu3zqtS2Btx
      cW6EGC3/xGDxI0BaU9Itr/73vwH3ZNilw2uE0xbhxkehShmL1hZns2A==

ARC-Message-Signature: i=1;a=rsa-sha256;c=relaxed/relaxed;d=securemx.jp;h=
      Content-Type:MIME-Version:Subject:Message-Id:To:From:Date:DKIM-Signature;s=
      arc20190303;t=1713333205;x=1713938005;bh=KXMY6AzIdALlOzgkBat5CPcMk0Vqq3IozV
      vMgC0v4D0=;b=heR13lSXC71wKTNdzf9UmnChz5/bqf5L2qZ5X7i8xTsk2x2o42rqKN9Fzpweeb
      Lplf9Msm09ki7XWRJMQci3oi+Ut5Z7Hmjo3UdTrbMjaT621GPKID+gVbDpKxCDlWiCzc5ox3wrw
      MYDoM6A/2rQwjXsgTMoWovEmxyB795Db1y+J5dyCzkpWcJ/2JEXBU0PbQ4Qr1Ofkzq7Rsum5Cyi
      lHeXJw64GDbbh6uY5j1HM92CtGo9Ejw+r6c63snjJURvfrjCcB11hgrx84CclJw9klt9PtUct7P
      +l90RENzLmpelNYa0w9Y2LAGjRztzayIunJnvw0WJ1unizóiuufwJkA==
```

Section 3.4.4 of RFC 6376, which ARC follows, states that trailing line breaks must not be removed even in relaxed mode.

M365, however, was calculating hash values with such trailing line breaks removed, which is why the bh tag values did not match.

When we contacted Microsoft citing this RFC specification, the company acknowledged that the issue was on the M365 side.

2.2.4 Conclusion

ARC adoption rates remain low compared with other sender authentication methods such as SPF, DKIM, and DMARC, with many email systems yet to implement ARC authentication. Possible reasons for this are that not setting up ARC authentication does not currently cause major problems, and that RFC 8617, the ARC RFC, is still designated as Experimental.

Google's sender guidelines also recommend using ARC authentication if regularly forwarding email. Additionally, if ARC is not implemented properly, this can prevent SPF and

DKIM verification from being performed correctly during the forwarding process, and depending on DMARC policies, emails may be rejected, potentially affecting email deliverability.

The cooperation of intermediary systems that forward email is crucial for ARC. IIJ will continue supporting sender authentication technologies, including ARC, and strive to improve email reliability and deliverability.

2.3 Sharp Surge in Phishing Emails Targeting Japan

At the end of 2024, we observed one of the largest volumes of phishing emails on record on IIJ's email services. Figure 2 plots the number of emails received by IIJ-operated honeypots that were classified as spam.

From around the end of November, the total volume of phishing emails increased sharply, peaking around year-end. The phishing emails during this period spoofed entities such as Amazon, Sagawa Express, tax offices (e-Tax), and the ETC usage inquiry service, with the content of all messages being fake replicas of legitimate emails. The tax office phishing emails, in particular, coincided with the tax filing season, presumably in a bid to increase attack success rates.

According to the Council of Anti-Phishing Japan's monthly report on phishing[*2], December 2024 saw the highest number of phishing reports on record, mirroring the trend IIJ had observed. These findings indicated that not only IIJ but other ISPs also were observing a similar trend, and that many phishing emails targeting Japan were being sent out. The report's summary notes that Japanese organizations are lagging behind when it comes to security measures such as sender authentication, making it easier for phishing emails to reach users than in other countries.

We continue to observe high levels of phishing emails in 2025, and continued vigilance is crucial.

Incidentally, we also reported on such a surge in phishing emails in IIR Vol. 51[*3], published in June 2021. That was back during the rapid rise in remote work amid the COVID-19 pandemic, and we also observed phishing emails exploiting that situation. And in 2019, Emotet, which spread via password-protected ZIP file attachments, wreaked havoc in Japan, with widespread impacts being confirmed.

Hence, phishing emails and viruses repeatedly cycle through periods of dormancy and resurgence, constantly changing tactics and adapting to the times.

Security measures are therefore not a set-and-forget affair; they must be continuously reviewed and strengthened in response to evolving threats. This is an endless battle and an ongoing investment, and it requires organization-wide effort. IIJ will continue its unwavering efforts to maintain a safe and secure environment.

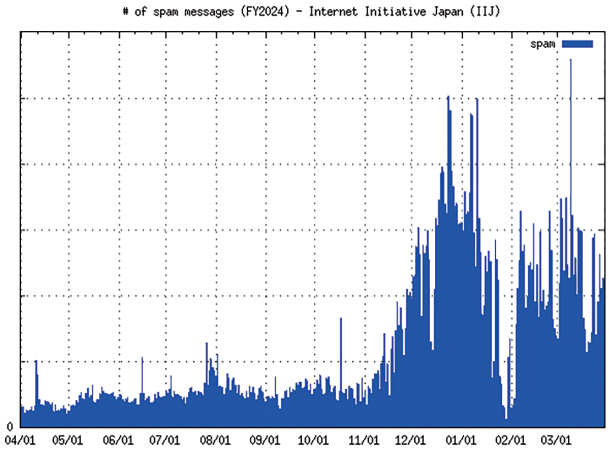

Figure 2: Spam Emails Arriving at IIJ Honeypots

*2 Council of Anti-Phishing Japan, Monthly Report: December 2024 Phishing Report Status (https://www.antiphishing.jp/report/monthly/202412.html, in Japanese).
*3 Internet Infrastructure Review (IIR) Vol.51 (https://www.iij.ad.jp/en/dev/iir/051.html).

2.4 Sender Authentication and Transport Encryption Statistics

Figures 3–6 show a percentage breakdown of sender authentication results aggregated from IIJ's email services as part of ongoing monitoring. The observations cover the month of March 2025.

Since our previous report, the percentage of successful sender authentications (pass) has decreased across all metrics. We know that sender authentication adoption in Japan, particularly DMARC, has increased significantly following the release of Google's sender guidelines[*4*5].

Thus, when this is considered alongside the data in Figure 2, the natural interpretation is that phishing emails that do not support sender authentication or that fail authentication checks have become dominant in the data relative to normal email, causing the overall pass rate to decline

Starting with this report, we have added aggregate results for ARC. Not all organizations are necessarily required to support ARC, but the IIJ Secure MX Service for enterprise email security added support for ARC signatures for received emails from 2019.

Next, we look at transport encryption (STARTTLS) data from the IIJ Secure MX Service. Figure 7 shows the types and percentages of transport encryption for received emails. PLAIN indicates no transport encryption.

Nearly 70% of received email communications used transport encryption, and TLSv1.3 accounted for close to half. As reported in Section 2.3, we saw a surge in phishing emails in December, and the graph here shows an increase in the proportion of TLSv1.2 around December. These results indicate that TLS-encrypted communications were also being used to send phishing emails at that time.

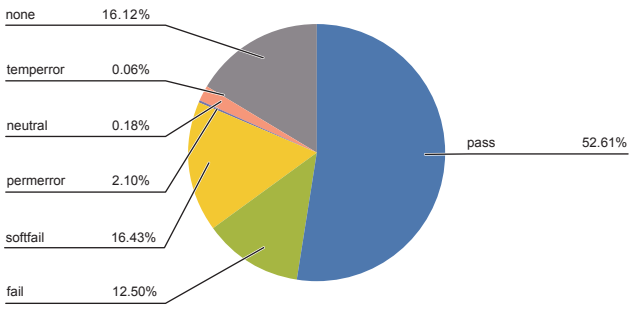

Figure 3: Breakdown of SPF Authentication Results

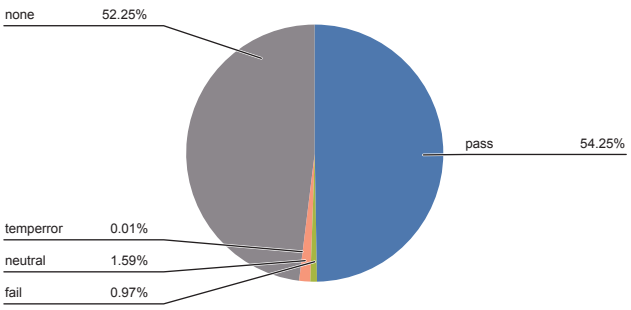

Figure 4: Breakdown of DKIM Authentication Results

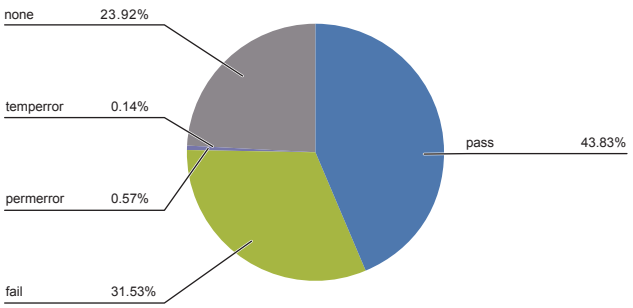

Figure 5: Breakdown of DMARC Authentication Results

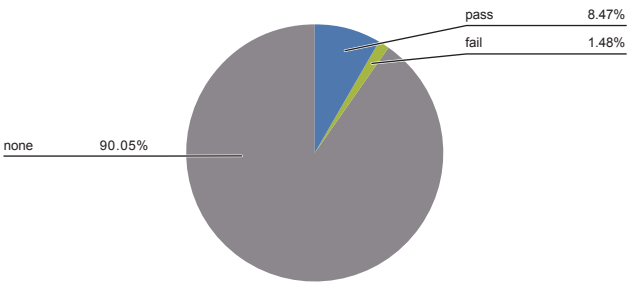

Figure 6: Breakdown of ARC Authentication Results

*4 Email sender guidelines, Google (https://support.google.com/a/answer/81126).

*5 DNSOPS.JP Statistics, Domain status of Japanese organizations – DNSSEC/SPF/DMARC (https://stats.dnsops.jp/chart/all/dmarc).

Figure 8 plots the proportion of emails sent from the IIJ Secure MX Service that used transport encryption. Due to equipment limitations, we can only show the aggregate proportion here, but it is evident that close to 100% of communications are encrypted.

When we last reported on this in IIR Vol. 59[*6], the proportion was fluctuating around 80–90%, so about two years on, it seems safe to say that email, like the Web, has entered the era of always-on TLS.

The Google Transparency Report[*7] also provides percentage data for email transport encryption, which indicate a largely identical trend. The inclusion of STARTTLS as a requirement in Google's sender guidelines released in 2023 likely had a substantial impact here.

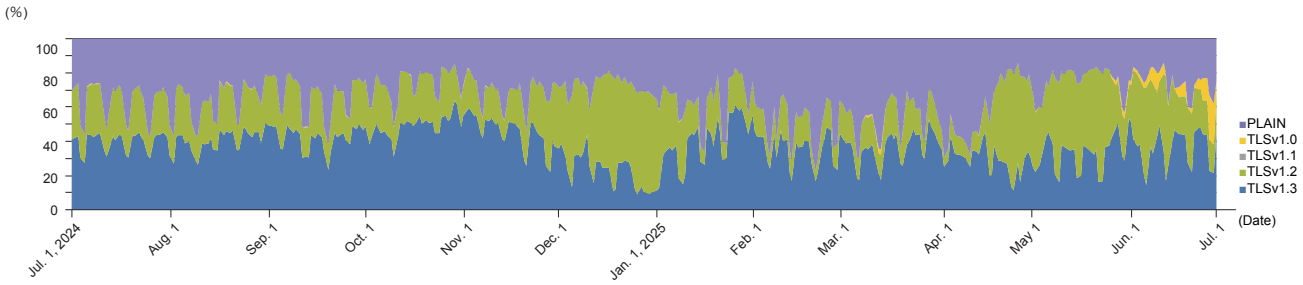

Figure 7: Proportion of Received Emails Using Transport Encryption

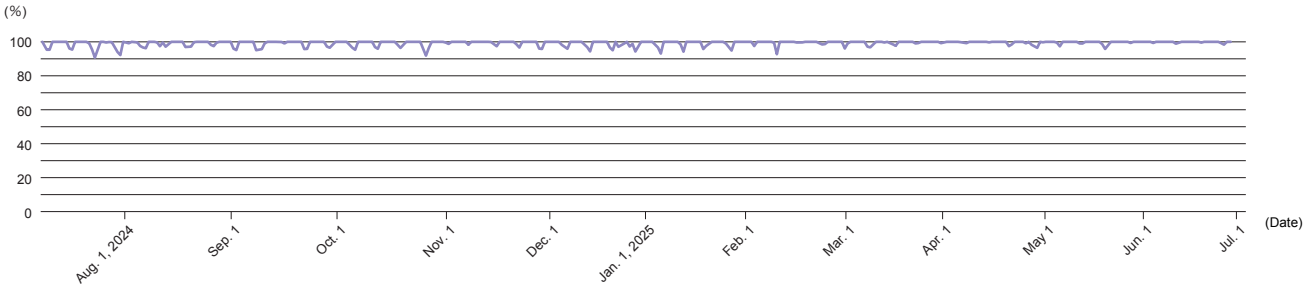

Figure 8: Proportion of Sent Emails Using Transport Encryption

2.1 New Initiatives to Protect Our Customers: Effectiveness of Defensive Action, 2.3 Sharp Surge in Phishing Emails Targeting Japan, 2.4 Sender Authentication and Transport Encryption Statistics

samu Koga

Manager, Application Service Management Section, Application Service Department 2, Network Division, IIJ
Mr. Koga joined IIJ in 2007. He is engaged in the operation of email services and ID governance management services. To keep customers' email boxes safe, he serves as a communicator and public speaker on the latest attack methods, trends in spam, and countermeasures. He is also involved in a wide range of community activities, including M3AAWG, WIDE Project, and openSUSE.

2.2 IIJ's Approach to Sender Authentication (ARC)

Shunpei Yamashita

Mail Service Development Section, Application Service Department 2, Network Division, IIJ
Mr. Yamashita joined IIJ in 2021. He is engaged in the development of email services.

*6 Internet Infrastructure Review (IIR) Vol.59 (https://www.iij.ad.jp/en/dev/iir/059.html).
*7 Transparency Report, Google (https://transparencyreport.google.com/safer-email/overview).

Answering Calls to Move Away from VMware —The Tight Relationship Between Kubernetes and Virtual Environments

3.1 The VMware Issue Shaking the Industry

The debate around calls to move away from VMware has been raging since last year. Some say VMware license prices have surged and organizations should consider switching to other virtualization platforms. Others say vendor logic is being given precedence while existing customers are left neglected. Is this really true?

While various conditions make it difficult to generalize, a dispassionate look across the entire virtualization market would indicate that the company's actions are aimed at positioning suite packages, which have become the main battleground, as VMware's core offering as well, in an attempt to guide customers toward higher value-added, more competitive packages. Depending on how you look at it, you could say it is now less strategically differentiated from competitors. Not to favor any specific competitor's products, but I think it's reasonable to describe this as rationalization, whether in terms of lineup or pricing. The fact remains, however, that users of simple packages, most likely those in the volume zones, have borne the brunt and been hit by cost increases.

This leads people to consider taking some form of action, but it's highly questionable whether simply focusing on license fees and considering a switch to alternative products would achieve the desired outcome. After all, it's not as if VMware's appeal as a product has diminished. It may be a different story if version upgrades had introduced unwanted feature changes or problems that hurt customer trust, but this is purely a cost issue. Even if options that offer advantages in terms of license fees can be found, they would still need to be on par with VMware, which remains the virtualization market leader, in terms of functionality, stability, future prospects, and many other factors. Plus, moving away from VMware involves significant migration costs. And such costs might even offset the expected cost reduction benefits.

In other words, to successfully move away from VMware, one needs to think about what sort of platform would be even better than VMware, rather than fixating on the costs. Achieving both could not only deal with the licensing issue but also bring a competitive advantage to the business by modernizing the platform.

One answer to this extremely difficult challenge that IIJ is working on is the migration from VMware to Kubernetes. By using open source to the maximum extent and positioning our internally developed Kubernetes distribution, IKE (IIJ Kubernetes Engine), as a VMware alternative, we are significantly reducing license costs while also achieving operational efficiency and quality improvements unattainable at the IaaS layer.

3.2 Kubernetes Swallowing Up All Kinds of Workloads

Even engineers fairly well versed in cloud-native technology mostly consider Kubernetes a container orchestrator, I think. Indeed, it surely is the case that the vast majority of Kubernetes deployments are being used for container orchestration. But are you aware that in recent years, Kubernetes is gradually being adopted more and more not just for containers but for VMs as well? If you do know this, you are probably already quite in the know. I consider myself to be a cloud-native advocate, and even I felt that, while the technology is maturing, it would still be some time before we'd see adoption in production environments.

And yet, amid all this debate sparked by the VMware situation about reconsidering what platforms are being used, the tide seems to have turned dramatically, with Kubernetes having suddenly emerged as the leading potential successor. I'm not only talking about the mood here at IIJ. This is a trend I sense across the entire IT industry.

The reasons Kubernetes has come to attract attention as a VMware alternative are too numerous to list. While there certainly are technical reasons, I feel reliability plays an even bigger role. For example, it's an open source project operated in a vendor-neutral manner by the Linux Foundation, and there is little concern about vendor lock-in. It's widely supported across the entire IT industry, by both hardware and software vendors, with a rapidly growing ecosystem. As a result, much like Linux in the server OS space, Kubernetes has come to be regarded as the de facto standard for orchestrators, something organizations can commit to for long-term investment. It is this confluence of factors that has earned it trust from users.

That said, if Kubernetes's role had remained limited to being a container orchestrator, it might never have been regarded as mainstream. While container adoption is increasing rapidly on server-side systems, if one asks whether VMs or containers are the primary tool, the answer, at present, is clearly VMs. But on the other hand, it seems unlikely that platforms only capable of managing VMs will continue to be the mainstream forever. Given its ability to handle VMs and containers as equivalent workloads, it seems inevitable that Kubernetes would attract attention.

3.3 IKE, IIJ's Kubernetes Platform

IIJ is gradually migrating from VMware to Kubernetes, and what enabled us to select Kubernetes as our means of moving away from VMware at such an early stage is that we already had sufficient operational experience with Kubernetes as a container platform. If the move away from VMware had spurred us to tackle Kubernetes operations for the first time, we may have hesitated quite a bit.

That's because Kubernetes operational knowhow isn't something you can acquire overnight. Kubernetes is often referred to as the OS for the cloud era since it mediates system resources like servers, networks, and storage. Moreover, unlike OSes that control resources within the confines of a single server, a single Kubernetes instance can manage all manner of large-scale systems housed in a data center so that they are coordinated with each other. To accomplish this, Kubernetes abstracts away the physical systems (hides the physical divisions between them), making the data center appear as if it were one giant resource pool. And users perform all operations by issuing commands through Kubernetes, just as they would with a public cloud. This is indeed one of the motivations for using Kubernetes, but when it comes to designing and operating such abstracted systems, even experienced engineers may find that their existing knowledge is not enough. On the other hand, becoming well-versed in Kubernetes alone only confers operational abilities, not the ability to manage systems. Even with Kubernetes, you still ultimately end up having to deal with hardware control, so unless you understand how Kubernetes operations are reflected in the actual systems, you may find it difficult to provide sufficient quality of operations.

While similar issues exist with virtualization platforms like VMware, Kubernetes is even more complex due to its higher level of abstraction.

IIJ, meanwhile, introduced Kubernetes as a service platform back in 2018. At the time, we were using Kubernetes v1.9, which was a far simpler system than what we have today. The 23 upgrades since then bring us up to v1.32. Over the intervening time, we have also upgraded or replaced many of the plugins and controllers we were initially running, leaving barely a trace of the components of the original Kubernetes cluster. Nevertheless, the many workloads that have been running on this Kubernetes cluster since the early days continue to operate stably.

This is an extremely important point. It means that, even while countless features have been added to Kubernetes with each minor version upgrade, not once in seven years has there been a fatal loss of compatibility. Moreover, the fact that the impact on systems running on Kubernetes has been minimal despite implementations, both hardware and software, having changed dramatically at the lower levels is evidence of the effectiveness of Kubernetes's abstraction. This speaks to Kubernetes's high level of continuity, stability, extensibility, and future potential, and we can no doubt expect that to continue ahead.

That said, this didn't happen by accident. It is the result of making upgrades in an appropriate manner based on a detailed understanding of the content of each upgrade and the scope of its impact. The experience gained through that process has been tremendously helpful in our move away from VMware. Of course, there are any number of ways to acquire Kubernetes operational skills in less time, but it still takes some time to gain confidence.

3.4 The Networking Challenges of VMs on Kubernetes

While we had accumulated the requisite knowledge, we actually only started deploying VMs on Kubernetes at IIJ a year ago. When we began our evaluation, we were considering separating our infrastructure into Kubernetes for containers and Kubernetes for VMs, but we quickly realized this sort of thing was completely unnecessary. Indeed, one could even say that a real advantage of Kubernetes is the ability to mix and match containers and VMs according to the use case and build systems that combine both. That's how compelling the environment made possible by KubeVirt is. At the risk of

oversimplifying, most Kubernetes functionality applies to VMs as well if you mentally substitute VirtualMachine (the resource used to launch VMs) for Pod (the resource used to launch containers).

Treating Pods and VirtualMachines as equivalent workloads provides substantial benefits.

- As Kubernetes evolves, you enjoy the benefits for both containers and VMs without doubling your effort.
- For on-premises Kubernetes, you can maintain high utilization rates since there's no need to split infrastructure for containers and VMs.
- Pods and VirtualMachines attach to the same pod network, making interoperability a breeze.
- Replacing systems built as VMs with containers is easy. This is useful as a migration path from VMs to containers

Yet Kubernetes as a path away from VMware is not all smooth sailing. In many existing Kubernetes environments, the common pattern is to attach all containers and VMs to a single pod network per cluster. By contrast, VM-centric environments often provision multiple L2 networks and assign independent networks for each application. While you can provision equivalent networks on Kubernetes, this is not yet common practice, and L2 implementations frequently depend on the underlying infrastructure or proprietary networking products, so careful evaluation is needed. Where complex network topologies are required, simply extending an existing Kubernetes cluster may not suffice; you may need a Kubernetes cluster dedicated to the VMware exit.

That said, it's encouraging that the substantive challenges are largely confined to networking. For users, the skills needed to use VMs are truly minimal. Only a small amount of training is needed to get started. While it's a different story on the operations side, if the burden falls mainly on the platform engineers, that's actually gratifying because it's exactly what we're here for.

Our move away from VMware is still a work in progress, but there's no doubt that Kubernetes can be one of the answers. That said, requirements obviously vary widely by workload, and it's equally true that there can be multiple right choices. To deliver value to our customers through our services, we aim to set aside preconceived notions and choose the best option for each case.

Keisuke Taguchi
General Manager, SRE Promotion Department, Network Division, Network Service Business Division, IIJ
Mr. Taguchi has been involved in the launch of numerous services including email, DNS, server hosting, and cloud IaaS services. In recent years, he has drawn on his past experience to establish the platform engineering department, nurturing it into a platform that hosts over 100 services and projects. He is a strategist who believes staying abreast of changes in markets and technology and continuously updating one's knowledge and skills is the key to business success.

The 3G Sunset and MVNOs

4.1 Introduction

NTT DOCOMO's 3G service will end on March 31, 2026. Details can be found at the links below.

- Official NTT DOCOMO information
Notice of termination of FOMA and i-mode services
https://www.docomo.ne.jp/info/3g_closed/index.html (in Japanese)

- Official IIJ information
Regarding continued use of IIJ Mobile following the shutdown of NTT DOCOMO FOMA (3G)
https://www.iij.ad.jp/svcsol/mobile-support/news/3g_closed.html (in Japanese)

Regarding the NTT DOCOMO FOMA (3G) shutdown (IIJmio Mobile Type D)
https://www.iijmio.jp/info/iij/1712655772.html (in Japanese)

Support for mobile data communication features related the end of NTT DOCOMO 3G service
https://support.seil.jp/ (in Japanese)
→ Support → Technical information → Support for mobile data communication features related the end of NTT DOCOMO 3G service

NTT DOCOMO announced information back in 2019 about the termination of service, and the final date next March is now approaching. This will render 3G services from MVNO operators including IIJ unavailable, affecting voice calls, SMS, and data communication.

Devices that only support 3G will naturally become unusable, but what's not widely known is that even 4G (LTE)-capable devices may become unusable as a result.

This article provides a technical explanation of why this will happen. With there being little time left, we hope this will help many people navigate through the end of March 2026 by preparing for the 3G shutdown and avoiding the sudden loss of access to communication services.

4.2 Why is the 3G shutdown necessary?

Before discussing the problems that will arise from the 3G sunset, we first look at the evolution of cellular mobile wireless technology. Today's mobile wireless technology has evolved with new wireless standards being introduced for commercial services roughly every 10 years since the introduction of 2G in 1991 (Figure 1).

Meanwhile, when new wireless standards are introduced:

- Investment in equipment for previous-generation wireless standards drops substantially, with the focus shifting to maintenance and upkeep of existing facilities only.
- As a result, equipment vendors stop manufacturing products and gradually discontinue support.
- Maintaining older equipment thus becomes increasingly difficult as the years pass.

This leads to the need to replace older wireless equipment with newer equipment. Further, compared with older wireless standards, newer standards define technologies that use the limited spectrum resources more efficiently, making migration to newer standards a natural progression.

At some point it thus becomes necessary to discontinue use of older wireless standards and migrate to newer ones. This is precisely why the 3G sunset is coming; it is a natural progression. As a supplement to Figure 1, we show the progression of 2G/3G shutdowns by carriers worldwide (Figure 2), as published by the GSA (Global mobile Suppliers Association), a global industry association of device vendors. These 2G/3G shutdowns are set to peak in 2025, with networks shifting en masse to 4G and later wireless generations (Figure 2).

In the early stages of a new wireless standard's introduction, however, the new wireless network starts off with limited coverage, so wireless networks are typically constructed so as to use the legacy standards to supplement coverage. To maintain continuity between wireless networks, the core

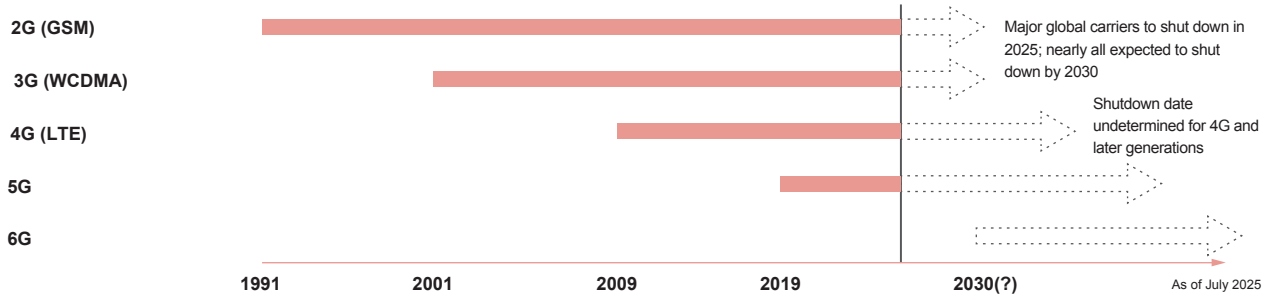

Figure 1: Global Service Launch Dates and Usage Periods for Cellular Mobile Wireless Technologies

networks work together to ensure users can move between the old and new wireless networks seamlessly without being disconnected.

With this in mind, devices are designed to support both old and new wireless standards, so that there are no problems when both old and new wireless networks exist alongside each other. But when old wireless networks are discontinued, as with the 3G sunset, this causes problems for devices designed on the premise that networks will be able to work together like this. The following explains this problem in detail.

4.3 Impact of the 3G Sunset

As mentioned in the previous section, device implementations mean that the 3G shutdown will have the following impacts.

(1) Devices only supporting 3G (no 4G support): Unusable as 3G signals are non-existent
(2) Devices support both 3G and 4G: May become unusable due to absence of 3G signals even when 4G signals are present

The problems with 3G/4G dual-compatible devices in (2) above can be further classified as follows.

(2-1) Problems caused by voice-capable devices (smart phones etc.)
(2-2) Problems caused by devices always connecting to a 3G network first

We explain each of these problems in the following section.

4.4 Problems Caused by Voice-Capable Devices (Smartphones etc.)

According to the 3GPP specifications defining mobile wireless standards, 4G devices can be broadly classified into two modes (UE usage settings). We explain each below.

(1) "Voice centric" mode

- A mode in which the device must connect to networks supporting voice functionality
  - The majority of smartphones use this implementation
  - Operates normally if either VoLTE (voice communication using 4G wireless) or CSFB (CS Fall Back), allowing voice communications by switching to 3G networks, is available

- However, when connecting to a 4G network, this mode causes the following problems when neither VoLTE nor CSFB is available.
  - Standard 3GPP behavior in this case is to disconnect

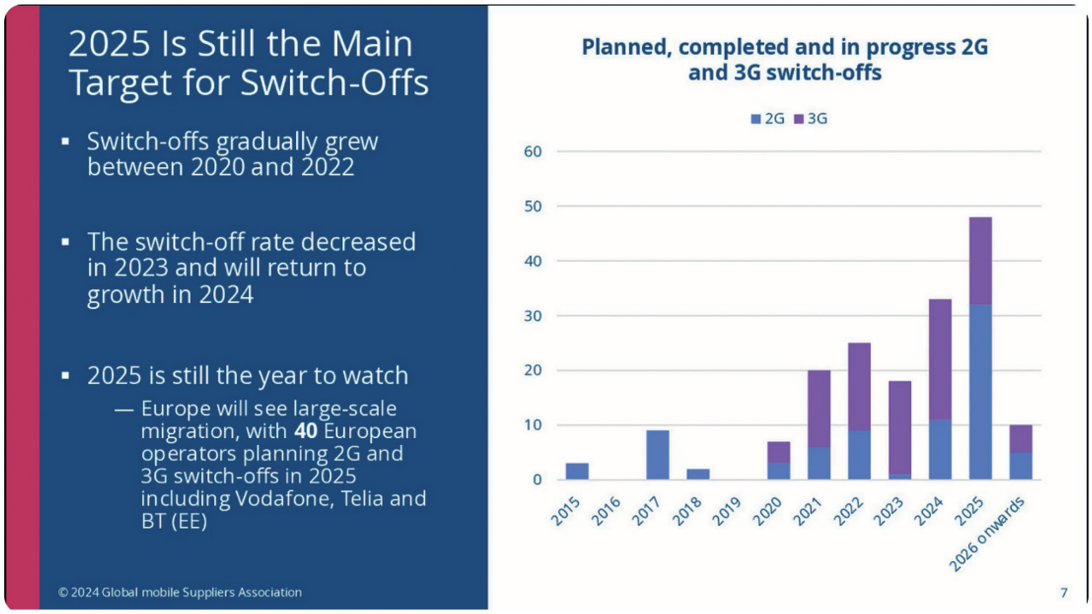

Source :https://gsacom.com/webinar/2g-3g-sunset-and-implications-for-5g-broadcast/

Figure 2: Courtesy of GSA, 2G/3G Sunset and Implications for 5G Broadcast

from the 4G network and switch to a mode that searches for another network (such as 3G) where voice is available
- If no available 3G network signals are found, the device remains out of service, so data cannot be sent or received
- After the 3G sunset, this problem may arise under certain combinations of device implementation and communication service constraints

- Cases that will be problematic after the 3G sunset include the following.
  - When using 4G (LTE) devices that don't support VoLTE
  - When using VoLTE-compatible 4G (LTE) devices on a service with no support for VoLTE

(2) "Data centric" mode
- The majority of data devices (routers, USB modems, communication modules) use this implementation
- Does not lose service like (1) even when voice functionality is unavailable
- Does not suffer from out-of-service issues after the 3G sunset

IIJ realizes that devices set to voice-centric mode (see (1)) may experience issues after the 3G shutdown, so for the Full MVNO-based IIJ Mobile Service/Type I, we plan to implement measures on the network side to prevent such issues.

4.5 Problems Caused by Devices Always Connecting to a 3G Network First

This problem may be common with 4G (LTE) devices not in voice-centric mode. The devices may be set up to first connect to a 3G instead of a 4G network because of implementation details or settings. In such cases, the devices often then switch to a 4G network after a given delay such that users may not be aware that the issue is there.

When both 3G and 4G networks are available, this only amounts to initially connecting to 3G networks and then switching to 4G networks after a while. However, after the 3G shutdown, the absence of 3G networks prevents connection to 4G networks (Figure 3).

Since the problem is on the device side, it cannot be addressed on the network end, so it's up to device users to do something about it. If users encounter symptoms like this, they need to act quickly. IIJ's knowledge of the problem indicates that the causes can be classified as follows.

(a) Devices designed to always connect to a 3G network first and for which it is not possible to apply a fix or change
- Such devices will be unusable after the 3G sunset and will thus need to be replaced

Examples:
The 510FU and 520BU (USB modems) previously sold for IIJ's enterprise mobile services fall into this category.

(b) Devices designed to always connect to a 3G network first but which can be modified via a firmware update
- While unusable post 3G sunset on older firmware,

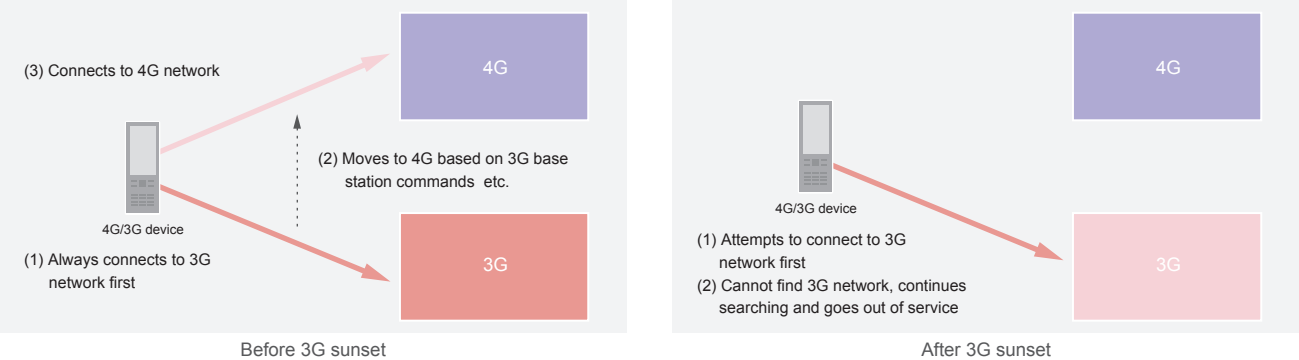

Figure 3: The Problem with Devices Connecting to a 3G Network First

such devices can be used after the 3G sunset if updating to the device manufacturer's latest firmware alters the device settings to first connect to a 4G network
- Whether a given model can be used after the 3G sunset must be confirmed with the manufacturer

Examples:
1. Fujisoft FS040U (secure connection mode)
2. NCXX UX302NC, UX302NC-R
3. Atmark Techno Armadillo-IoT G4/G3M1

(c) Devices that connect to a 3G network first by default but which can be configured to connect to a 4G network first
- Such devices will be unusable post 3G sunset if no action is taken but can continue to be used if the settings are changed
- Whether a device has such functionality and how to configure it must be confirmed with the manufacturer
- Key devices, and the settings required to force them to first connect to a 4G network, as follows.

Examples:
Google Pixel series (APN type needs to be set to "ia" in the APN settings)
Microsoft Surface series (Internet and attach settings APN setting need to be configured)

(d) Smartphones or communication modules for which the APN settings result in the device connecting to a 3G network first
- This problem is frequently reported by enterprise IIJ mobile service customers
- IIJ mobile service specifications require username, password, and authentication method to be entered in the APN settings and do not allow connections to be made without this information
- But for various reasons, values for username, password, and authentication method are not provided for 4G connections in certain cases, causing IIJ's equipment to reject such connections, thus precluding the devices from connecting to a 4G network
- After a 4G network connection is rejected in this manner, such devices fall back to 3G networks, and since the username, password, and authentication method are set properly for 3G network connections, such devices are then able to connect normally
- We omit the details here due to space constraints, but the reason this occurs has to do with differences between the initial connection method used on 4G and 3G networks
- Key devices and the associated fixes are as follows.

Examples:
Smartphones (all Apple devices)
The problem can be avoided by using iOS APN configuration profiles
Communication modules in general (Quectel EC25-J etc.)
The problem can be avoided by setting username, password, and authentication method in the communication module in advance via implementation-dependent AT commands (AT+QICSGP for Quectel)

4.6 Conclusion

This article has discussed how NTT Docomo's termination of 3G services on March 31, 2026 will also prevent certain 4G (LTE)-capable devices from connecting to 4G wireless networks. With less than a year until the 3G sunset, time is limited. The issues described in the latter half of this article often occur without users realizing it, so we are hopeful that readers will take this opportunity to double check the adequacy of their 3G shutdown preparations to ensure they can navigate smoothly through the 3G sunset.

Munenori Ohuchi
Senior Engineer, Platform Development Section, Infrastructure Development Department, MVNO Department, Mobile Services Division, IIJ
Mr. Ohuchi investigates and engages in research on cutting-edge mobile technology, and develops services utilizing such technology.

November 2025 Vol.67

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-MKTG020-0065

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/