IIR

Internet
Infrastructure
Review

May.2023
Vol. 58

Periodic Observation Report

SOC Report

Focused Research (1)

Interplay Between Data Centers and Electricity Markets

Focused Research (2)

The New IIJ Studio TOKYO’s Bridge to the Future

Focused Research (3)

The IIJ Backbone—30 Years of Transformations

IIJ
Internet Initiative Japan

Internet Infrastructure Review

May 2023 Vol.58

Executive Summary

ChatGPT has caused major waves since its release by OpenAI in November 2022. As its name suggests, ChatGPT is a type of chatbot. Chatbots have been used in all sorts of scenarios over the years. ChatGPT, however, has more advanced capabilities than conventional chatbots. For instance, it can respond naturally to questions in a wide range of fields, remember past conversations, and compose well-formed prose and computer programs.

When I tried it myself, I was nothing but astounded by these features. We often start by consulting a search engine when investigating something we have little knowledge of, but I have a feeling that one might obtain meaningful results even more efficiently by asking ChatGPT and then digging in deeper and asking it more questions based on the answers it provides. Some point out that ChatGPT will quite nonchalantly provide incorrect information at times, but the same can be said of search engines, so I think that, as always, we still need to be able to assess the veracity of the information we receive. That much remains the same.

We have for many years looked to search engines to sift through the vast seas of information on the Internet and make it easier for us to get at what we need. AI has played a part behind the scenes, but it is quite refreshing that we now have an AI chatbot as the actual user interface, and that it is even capable of summarizing information when needed. With search engines that provide a chat interface using OpenAI's GTP-4 already appearing, I sense real potential for change in the long-entrenched user experience that has so far characterized the search for information on the Internet.

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.

Chapter 1 presents our SOC Report, our periodic observation report for this edition. IIJ's SOC analyzes data obtained through the operation of IIJ's services, data that it collects independently, and data from external sources. Our report in this edition looks back on major security topics in 2022 and discusses those that were of particular interest to IIJ's security analysts, namely Emotet, VPN device vulnerabilities, and four other vulnerabilities frequently observed by our SOC in 2022.

The focused research report in Chapter 2 looks at electric power, something that is indispensable to the information & communications industry. As data transfer and compute volumes rise precipitously alongside the advance of information & communications technology, we naturally face a very real need from an environmental preservation perspective to keep the amount of power the industry consumes in check. The report here discusses challenges facing the electricity market in Japan and IIJ's efforts as a data center operator that consumes electric power.

The focused research report in Chapter 3 introduces IIJ Studio TOKYO, IIJ's video streaming center, and goes over the technologies involved. As networks and devices have evolved, video streaming has become part of our world in all sorts of areas. Internet-based video streaming represents a convergence of the video and ICT industries. I am pleased to present this report on the challenges IIJ has tackled in video distribution, including our efforts at IIJ Studio TOKYO.

And to coincide with IIJ's 30th anniversary, Chapter 4 presents a focused research report on the way the IIJ backbone network, which sits at the heart of our business, has transformed over the years. Since IIJ was founded, the Internet has continued to expand, and not only has IIJ's network grown but demands for availability, quality, and security have also risen incredibly. We wanted to use this edition of the IIR to chronicle how we have evolved our network over the years to adapt to these changes.

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.

SOC Report

1.1 Introduction

IIJ maintains the wizSafe security brand and works constantly to create a world in which its customers can use the Internet safely. The SOC communicates a variety of information on security issues via the wizSafe Security Signal[*1] site and conducts analyses of threat information using IIJ's Data Analytics Platform, which collects logs from IIJ services.

Section 1.2 of this report looks back at major security topics in 2022, and Section 1.3 discusses threats related to the topics covered, with a focus on those observed on our data analytics platform.

1.2 2022 Security Summary

Tables 1 and 2 show the security incidents that the SOC focused on from among those that rose to prominence in 2022.

[*1] wizSafe Security Signal (https://wizsafe.iij.ad.jp/).

Table 1: Incident Calendar (January–May)

Month	Summary
January	**A software developer announced it had been compromised by ransomware in the early hours of December 31, 2021. It was found not only that files had been encrypted in the attack but also that the attacker had stolen customer transaction and other data, which was published twice on data leak sites.** (Tokyo Computer Service Co., Ltd.) https://www.to-kon.co.jp/ja/topics/topics20220411103030/main/0/teaserItems1/0/linkList/0/link/letter.pdf https://www.to-kon.co.jp/ja/topics/topics20220411103030/main/0/teaserItems1/0/linkList/00/link/letter2.pdf https://www.to-kon.co.jp/ja/topics/topics20220411103030/main/0/teaserItems1/0/linkList/01/link/letter3.pdf https://www.to-kon.co.jp/ja/topics/topics20220411103030/main/0/teaserItems1/0/linkList/02/link/letter4.pdf
January	**A payment service provider announced that it had been subject to unauthorized access between August 2, 2021 and January 25, 2022. The attacks were complex and included unauthorized logins to internal management systems, SQL injection attacks into some applications, and the installation of backdoors, and it was revealed that personal information had been leaked.** (Metaps Payment Inc.) https://www.metaps-payment.com/company/20220125.html https://www.metaps-payment.com/company/20220228.html
February	**The JPCERT Coordination Center (JPCERT/CC) issued an alert on Emotet infections spreading rapidly since February. In March, the number of Emotet-compromised .jp email addresses that could be exploited to send emails surged to over five-fold the 2020 wave peak, and while these observations eased off in mid July, mailouts were again observed from November 2.** (JPCERT/CC) https://www.jpcert.or.jp/english/at/2022/at220006.html
March	**A major automaker announced it had suspended operations at all of its domestic plants (28 lines across 14 plants) on March 1. This was attributed to a system failure at a domestic supplier, which disclosed that this involved file server virus infections.** (Kojima Industries Corporation) https://www.kojima-tns.co.jp/wp-content/uploads/2022/03/20220331_%E3%82%B7%E3%82%B9%E3%83%86%E3%83%A0%E9%9A%9C%E5%AE%B3%E8%AA%BF%E6%9F%BB%E5%A0%B1%E5%91%8A%E6%9B%B8%EF%BC%88%E7%AC%AC1%E5%A0%B1%EF%BC%89.pdf https://www.kojima-tns.co.jp/wp-content/uploads/2022/08/%E3%82%A6%E3%82%A3%E3%83%AB%E3%82%B9%E6%84%9F%E6%9F%93%E8%A2%AB%E5%AE%B3%E3%81%AB%E3%82%88%E3%82%8B%E3%82%B7%E3%82%B9%E3%83%86%E3%83%A0%E5%81%9C%E6%AD%A2%E4%BA%8B%E6%A1%88%E7%99%BA%E7%94%9F%E3%81%AE%E3%81%8A%E7%9F%A5%E3%82%89%E3%81%9B-2.pdf
March	**An auto parts manufacturer announced on March 10 that a third party had illegally accessed the network at one of its overseas group companies.** (Denso Corporation) https://www.denso.com/global/en/news/newsroom/2022/20220314-g01/
March	**VMware announced on March 31 that the Spring Framework had a vulnerability (CVE-2022-22965) that was leaked out ahead of the CVE being published. If exploited, the vulnerability could allow remote arbitrary code execution. In addition to VMware products, this also impacts other products using the Spring Framework. The vulnerability is known as Spring4Shell.** (VMware) https://spring.io/blog/2022/03/31/spring-framework-rce-early-announcement
April	**VMware announced on April 6 that multiple vulnerabilities were present in some of its products, including VMware Workspace ONE Access and VMware Identity Manager. VMware also confirmed that, of these, a remote code execution vulnerability (CVE-2022-22954) and a local privilege escalation vulnerability (CVE-2022-22960) were being exploited in the wild.** (VMware) https://www.vmware.com/security/advisories/VMSA-2022-0011.html
April	**An industrial parts manufacturer announced on April 8 that a third party had illegally accessed the network at an overseas plant of one of its group companies. It reported that the intruder had likely exploited a VPN device vulnerability.** (Nippo Ltd.) https://www.nip.co.jp/news/.assets/20220408-1.pdf https://www.nip.co.jp/news/.assets/20220422-1.pdf
May	**On May 4, F5 Networks disclosed a vulnerability (CVE-2022-1388) in F5 BIG-IP that may allow iControl REST authentication to be bypassed. By exploiting this vulnerability, an unauthenticated attacker could execute arbitrary system commands, create or delete files, or disable services.** (F5 Networks) https://support.f5.com/csp/article/K23605346
May	**A cloud service provider announced that unauthorized access had been gained to some of the load balancers provided on its services between May 7 and 11.** (Fujitsu Cloud Technologies Limited) https://pfs.nifcloud.com/cs/catalog/cloud_news/catalog_202205161000_1.htm https://pfs.nifcloud.com/cs/catalog/cloud_news/catalog_202205311000_1.htm https://pfs.nifcloud.com/cs/catalog/cloud_news/catalog_202206071000_1.htm https://pfs.nifcloud.com/cs/catalog/cloud_news/catalog_202206291000_1.htm
May	**Microsoft announced on May 30 that a remote code execution vulnerability (CVE-2022-30190) was present in Microsoft Support Diagnostic Tool (MSDT). This zero-day vulnerability is known as Follina.** (Microsoft) https://msrc-blog.microsoft.com/2022/05/30/guidance-for-cve-2022-30190-microsoft-support-diagnostic-tool-vulnerability/

Table 2: Incident Calendar (June–December)

Month	Summary
June	**Atlassian announced on June 2 that an unauthenticated remote code execution vulnerability (CVE-2022-26134) was present in its Confluence Server and Confluence Data Center products.** (Atlassian) https://confluence.atlassian.com/doc/confluence-security-advisory-2022-06-02-1130377146.html
June	**A local government announced that a USB flash drive containing personal information had been lost on June 21. The USB flash drive was found on June 24, and an investigative report released on November 28 said there was no evidence any personal information had been leaked.** (City of Amagasaki) https://www.city.amagasaki.hyogo.jp/kurashi/seikatusien/1027475/1030947.html
July	**A hotel chain announced on July 27 that 11,961 personal information records had been breached owing to Emotet infecting a computer at one of its outsourcing contractors.** contractors. (APA Group) https://www.apa.co.jp/newsrelease/164149
September	**Access was disabled to 23 websites of four Japanese ministries/agencies (Digital Agency, Ministry of Internal Affairs and Communications, Ministry of Education, Culture, Sports, Science and Technology, and Imperial Household Agency) as well as websites of a number of private-sector companies. The attack is believed to have been carried out by the pro-Russia hacking group Killnet, which posted a video to its Telegram channel declaring war against the Japanese government.**
October	**On October 10, Fortinet disclosed a vulnerability (CVE-2022-40684) in the administrative interface of its FortiOS, FortiProxy, and FortiSwitchManager products. By exploiting this vulnerability, an attacker could bypass authentication and perform arbitrary operations on the interface.** (Fortinet) https://www.fortiguard.com/psirt/FG-IR-22-377
October	**It was disclosed on October 19 that over 65,000 company data records had been potentially exposed by the misconfiguration of an object storage service provided by a software developer. After a foreign security firm notified Microsoft of the issue, the misconfiguration was rectified and affected customers were notified.** (Microsoft) https://msrc-blog.microsoft.com/2022/10/19/investigation-regarding-misconfigured-microsoft-storage-location-2/
October	**The National Institute of Information and Communications Technology (NICT) announced that DDoS bot infections targeting DVRs manufactured by Focus H&S were increasing, with 17,489 attacks observed between June 1 and August 31. As the products are also sold in Japan, NICT is calling on users to update the firmware as soon as possible.** (National Institute of Information and Communications Technology) https://blog.nicter.jp/2022/10/analysis-of-ddos-bot-targeting-dvrs/
October	**A medical center disclosed on October 31 that it had received what was apparently a ransomware attack. The attack caused a failure in the center's electronic medical record system, making it impossible to perform regular consultations and also causing the suspension of general outpatient services.** (Osaka General Medical Center) https://www.gh.opho.jp/pdf/obstacle20221031.pdf
November	**A PC peripherals manufacturer disclosed on November 21 that a third party had gained unauthorized access to a website it operates, and that tampering with its payment application had resulted in the personal information of up to 147,545 people and 1,938 credit card records being breached.** (Wacom) https://www.wacom.com/ja-jp/about-wacom/news-and-events/2022/1484
December	**On December 12, Fortinet disclosed a vulnerability (CVE-2022-42475) in FortiOS SSL-VPN that, if exploited, could allow a remote unauthenticated attacker to execute arbitrary code or commands.** (Fortinet) https://www.fortiguard.com/psirt/FG-IR-22-398
December	**On December 14, Citrix disclosed a vulnerability (CVE-2022-27518) in Citrix Gateway and Citrix ADC that, if exploited, could allow an unauthenticated remote attacker to execute arbitrary code.** (Citrix) https://support.citrix.com/article/CTX474995/citrix-adc-and-citrix-gateway-security-bulletin-for-cve202227518
December	**A password management software provider disclosed on December 22 that an unauthorized party had gained access to the cloud storage it uses to store production data backups. The attack resulted in customer information and encrypted sensitive information, including information generated when customers use the service, being breached.** (LastPass) https://blog.lastpass.com/2022/12/notice-of-recent-security-incident/

1.3 Security Topics

This section discusses the SOC's observations on the key topics our analysts were focused on from among the 2022 security incidents covered in Section 1.2.

1.3.1 Emotet in 2022

■ Emotet Overview

Emotet activity rose to the fore again in 2022, after a quiet period in 2021 following the high volume of observations back in 2019 and 2020. Emotet was first reported in 2014, and at the time was known as a type of malware called a banking trojan, which steals financial details and the like. Added functionality subsequently saw Emotet transform, and it increasingly began to infect other devices. Starting with its observations in September 2019, in 2019 and 2020 our SOC observed many Emotet-related emails and C&C server communications. With Europol taking down Emotet's attack infrastructure on January 27, 2021[*2], Emotet infections disappeared for a time from January 26, 2021. But on November 14, 2021, Emotet was able to compromise systems via another malware named Trickbot[*3], and Emotet file downloads via email were detected. In 2022, the JPCERT Coordination Center (JPCERT/CC) reported that the Emotet infections were greater in scale than those in 2020, and as illustrated by examples given in the previous section's incident calendar, this rampancy did lead to some damage (Tables 1 and 2).

Emotet has functionality for stealing information and spreading itself, as well as loader and botnet functionality allowing it to download and execute other malware. Once it infiltrates a device, Emotet steals information like email addresses, account information, and email text. It also creates emails using stolen email text and subject lines, to which it attaches a file that downloads Emotet before sending out. It is known that it attaches VBA macro-containing Microsoft Office Word/Excel files and password-protected ZIP files that contain such Word/Excel files in compressed form. The password-protected ZIP files are encrypted, so the contents cannot be examined if they cannot be decrypted by security products such as antivirus or sandbox systems. Files that cannot be inspected can circumvent security products, so there is a good probability of them making it to the end user. A new approach was also observed in April 2022, whereby Emotet attaches a shortcut file (LNK file) or a password-protected ZIP containing a shortcut[*4], so care must be taken with other file formats, not just the conventionally used Word and Excel formats.

And Emotet's loader functionality, mentioned above, means it can also serve as an entry point for other malware. In 2022, Palo Alto Networks confirmed that devices compromised by Emotet also contained malware known as IcedID and Bumblebee[*5]. Cybereason also reported on attacks in which Emotet was used to deploy the Cobalt Strike framework, which is used in penetration testing[*6]. So if Emotet infections are left unaddressed, the damage caused can become even more serious. These attributes make Emotet one of the more infectious and threatening malware programs out there.

[*2] Europol (European Union Agency for Law Enforcement Cooperation), "World's most dangerous malware EMOTET disrupted through global action" (https://www.europol.europa.eu/media-press/newsroom/news/world%e2%80%99s-most-dangerous-malware-emotet-disrupted-through-global-action).

[*3] Cyber.wtf, "Guess who's back" (https://cyber.wtf/2021/11/15/guess-whos-back/).

[*4] Information-technology Promotion Agency (IPA), "Emails designed to cause Emotet malware infections" (https://www.ipa.go.jp/security/announce/20191202.html#L20, in Japanese).

[*5] Twitter (@Unit42_Intel) (https://twitter.com/Unit42_Intel/status/1590002190298804225).

[*6] Cybereason, "Threat analysis report: All roads lead to Cobalt Strike—IcedID, Emotet, QBot" (https://www.cybereason.co.jp/blog/malware/7797/).

Emotet Observations

Here, we report on Emotet observations by our SOC.

Figure 1 shows the number of Emotet emails over the course of a year. Date is on the horizontal axis. The detection count is on the vertical axis, normalized so that the total number of detections over the period corresponds to 100%.

Broadly speaking, Emotet was detected during three periods. The first was January 21 – April 5, which saw the highest detection count. Detections increased from February 2 and peaked on March 2. This corresponds to when JPCERT/CC reported a rapid increase in the number of Japanese email addresses possibly compromised and being abused by Emotet to send emails (Tables 1 and 2), and our SOC observed a similar trend.

The second period was April 21 – July 14, with observations increasing from June 4 and peaking on June 14. In addition to macro-bearing Excel files, detections from April 30 showed that Emotet was also attaching shortcuts and ZIP files containing shortcuts.

The third period was November 2–12, which saw the fewest detections. The shortcut files detected in the second period were not detected during this period; only macro-bearing Excel files were detected.

Even during the periods when attack emails were not observed (April and July–November), it was reported that Emotet itself had been changed[*7], so it would seem that activity is ongoing even during times when conspicuous attacks are not being observed.

Next, Figure 2 shows Emotet device infections on each date as a percentage of the year's total. Date is on the horizontal axis. The number of devices infected by Emotet is on the vertical axis, normalized so that the total

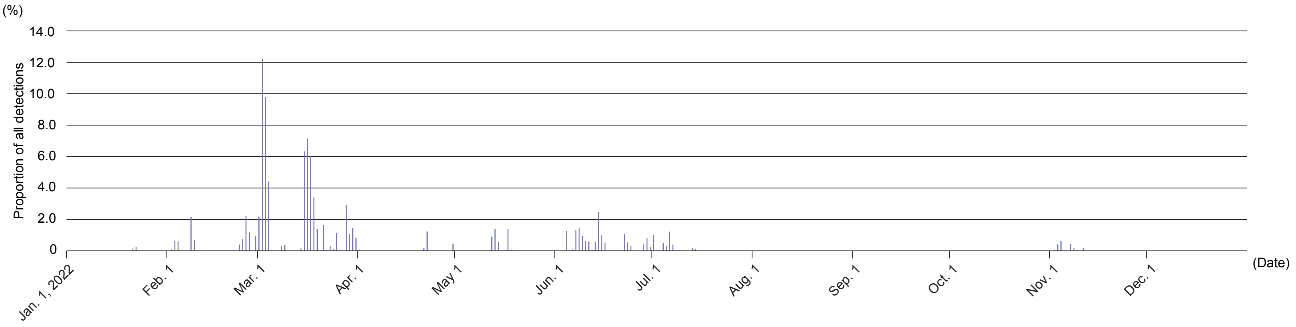

Figure 1: Number of Emails in which Emotet was Detected

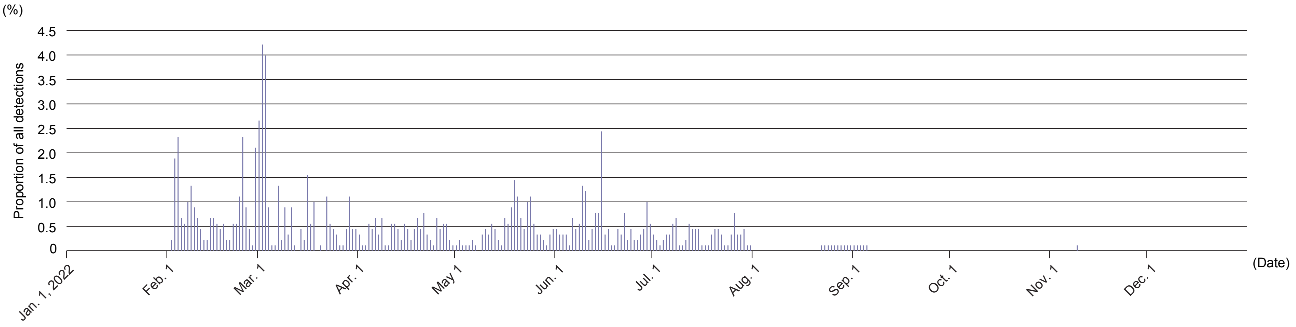

Figure 2: Devices infected by Emotet as Proportion of Annual Total

[*7] Proofpoint, "A Comprehensive Look at Emotet Malware's Fall 2022 Return" (https://www.proofpoint.com/au/blog/threat-insight/comprehensive-look-emotets-fall-2022-return).

corresponds to 100%. Each IP address that engages in communications with a C&C server is counted as one device.

As with the Emotet emails, the observations began on February 2 and peaked on March 2. From April to May, the number of devices infected tended to be lower than in March, but a large number of devices was again observed on June 15.

Periods of high device observations coincide with periods when a high number of emails in which Emotet was detected were received, but there are also periods in which infection traffic continued to appear despite no Emotet emails being received. This is probably because Emotet infections went unnoticed for some time after occurring, which suggests insufficient early detection measures. Emotet has loader functionality, and failure to detect an infection can result in other malware being let in, possibly resulting in serious damage, so you have to take steps to detect it early.

■ Countermeasures

Emotet infection activity has stopped as of end-December 2022, but it could resume, so Emotet is one form of malware that we should continue to keep a close eye on. Emotet infiltrates systems when users open VBA macro-containing Microsoft Office files and shortcut files (LNK files). A method of limiting the damage caused by infections is to disable the automatic execution of macros so that macros are not executed when a file is opened. Both the Microsoft Office files and the shortcut files execute PowerShell to download Emotet, so disabling PowerShell from running, if you are not using it in your operations, is one possible countermeasure. If possible, ending the practice of using password-protected ZIP archives to share files and changing your operations to block password-protected ZIP files across the board when emails are received is also an effective step. To spread itself, Emotet will also regurgitate actual legitimate emails, so it can be difficult for humans to spot Emotet-related emails. If you suspect that you have been infected with Emotet by executing a macro or following a shortcut, you can use EmoCheck, a tool released by JPCERT/CC, to check for this[*8]. Making sure you are set up to enable early detection and a rapid initial response is also key, and this can be achieved through monitoring by a SOC or EDR and installing antivirus software.

1.3.2 VPN Device Vulnerabilities that Expose Corporate Networks

Stories of major corporations, hospitals, and other Japanese entities suffering damage from ransomware infections appeared frequently in the news in 2022 (Tables 1 and 2). Businesses had to temporarily shut down in some cases depending on the damage caused, resulting in product deliveries being delayed, hospitals pausing patient intake, and so forth, so the impact of infections rippled through to customers and local communities.

The main triggers of ransomware damage are systems being infected by malware from suspicious emails or suspicious sites and attackers infiltrating organizations' networks. The rise of telework in recent years has led an increasing number of companies to install VPN devices so that workers can connect to internal networks via the Internet, and so instances of attackers exploiting VPN device vulnerabilities to infiltrate networks are becoming quite prevalent. According to a survey by Japan's National Police Agency, 68% of companies that suffered ransomware damage in the first half of 2022 cited VPN devices as the intrusion route[*9].

In this section, we look at known VPN device vulnerabilities and discuss our SOC's observations of attacks targeting those vulnerabilities.

[*8] GitHub (@JPCERTCC) (https://github.com/JPCERTCC/EmoCheck).

[*9] National Policy Agency, "Information on cyberspace threats in 1H of 2022" (https://www.npa.go.jp/news/release/2022/20220914001.html, in Japanese).

■ VPN Device Vulnerabilities

Table 3 lists 14 VPN device-related vulnerabilities mentioned in alerts[*10] issued by JPCERT/CC during 2018–2022. “CVE ID” is the Common Vulnerabilities and Exposures Identifier managed by the US-based MITRE Corporation, and “CVSS v3 base score” is a gauge of the severity of the vulnerability, ranging from 0.0 to 10.0, as assessed by the US National Institute of Standards and Technology (NIST) under the Common Vulnerability Scoring System.

■ SOC’s Observations of Attacks Targeting the Vulnerabilities

Attack activity aimed at discovering and exploiting vulnerable VPN devices is relentless on the Internet. Figure 3 graphs the number (as a percentage) of attacks targeting the VPN device vulnerabilities shown in Table 3 observed by our SOC in 2022. The data are normalized so that the total number of attacks detected during the period sums to 100% on the vertical axis.

Table 3: Main VPN Device-Related Vulnerabilities (2018–2022)

Date published	Product vendor	CVE ID	CVSS v3 base score	When exploitable	Impact
Dec. 2022	Citrix	CVE-2022-27518	9.8	Before authentication	Execute arbitrary code
Dec. 2022	Fortinet	CVE-2022-42475	9.8	Before authentication	Execute arbitrary code
Oct. 2022	Fortinet	CVE-2022-40684	9.8	Before authentication	Bypass authentication on admin interface
Dec. 2021	SonicWall	CVE-2021-20038	9.8	Before authentication	Execute arbitrary code
Sep. 2021	SonicWall	CVE-2021-20034	9.1	Before authentication	Delete arbitrary file
Apr. 2021	Pulse Secure	CVE-2021-22893	10.0	Before authentication	Execute arbitrary code
Mar. 2021	F5 Networks	CVE-2021-22986	9.8	Before authentication	Execute arbitrary code
Feb. 2021	SonicWall	CVE-2021-20016	9.8	Before authentication	Steal credentials and session information
Jul. 2021	F5 Networks	CVE-2020-5902	9.8	Before authentication	Execute arbitrary code
Dec. 2019	Citrix	CVE-2019-19781	9.8	Before authentication	Execute arbitrary code
Jul. 2019	Palo Alto Networks	CVE-2019-1579	8.1	Before authentication	Execute arbitrary code
May 2019	Fortinet	CVE-2018-13379	9.8	Before authentication	Download arbitrary files
May 2019	Pulse Secure	CVE-2019-11510	10.0	Before authentication	Download arbitrary files
Apr. 2019	Pulse Secure	CVE-2019-11539	7.2	Before authentication	Execute arbitrary code

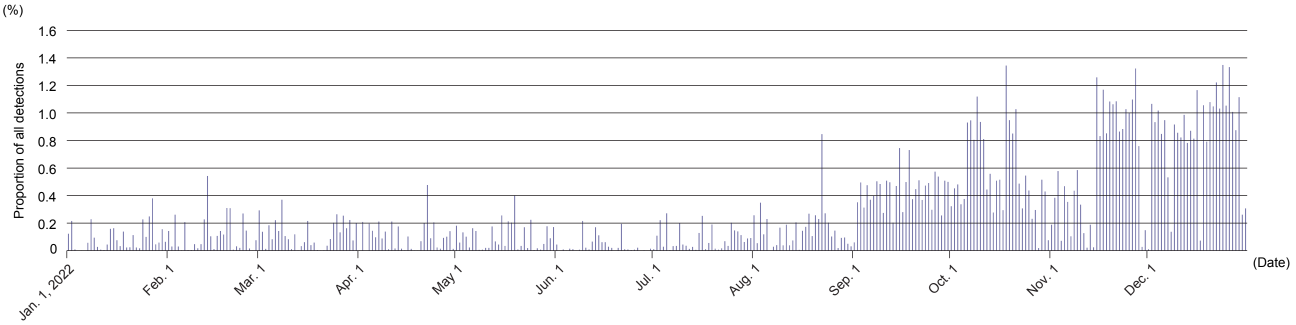

Figure 3: Observations of Attacks Targeting VPN Device Vulnerabilities (January – December 2022)

[*10] JPCERT/CC (https://www.jpcert.or.jp/english/at/2022.html).

As Figure 3 shows, attack activities targeting VPN device vulnerabilities persisted throughout the year, with the number of detections per day rising from September 2022. In Figure 4, we give a breakdown of detection counts (as a proportion of total) in 2022 for each of the CVE IDs in Table 3.

Of the vulnerabilities shown in Figure 4, our SOC analysts paid particular attention to the following.

CVE-2018-13379 (File-Stealing Vulnerability in Fortinet VPN devices)

CVE-2018-13379 was the most commonly detected vulnerability at 91.70% of total. Attackers can exploit this vulnerability to steal credentials needed to connect to VPNs. Despite more than three years having passed since it was published, many attacks were still being carried out, with damage actually being caused in cases. A hospital that had been the subject of a ransomware attack in October 2021 published a report[*11] in June 2022 saying that the attacker had likely exploited CVE-2018-13379 to infiltrate its systems. The medical center hit by a ransomware attack in October per the incident calendar (Tables 1 and 2) may also have been affected, as it was shown[*12] that this vulnerability was present in the OS version used on VPN devices of the medical center's trading partner identified as the likely intrusion vector.

CVE-2022-40684 (Authentication Bypass Vulnerability in Fortinet VPN devices)

CVE-2022-40684 is a relatively new vulnerability published in October 2022. Attackers can exploit this vulnerability to bypass authentication and access the administration interface of the VPN device, which could provide a foothold for infiltrating an organization's internal network by changing settings. Figure 4 shows this accounted for 1.48% of detections in 2022, putting it in the No. 4 spot, but most of these detections occurred over a mere 11 days, December 21–31, 2022. Ongoing vigilance is required as such attacks may continue in 2023 and beyond.

Preventing Your Organization from Becoming a Victim

Steps you can take to prevent your organization from becoming another victim include: updating the OS and firmware of your VPN devices, turning off unnecessary functionality, not externally exposing functionality that does not need to be (e.g., administration interfaces), setting source IP address restrictions to allow VPN connections, and setting up two-factor authentication.

It must be remembered that even if you have not yet discovered an intrusion into your internal network, an attacker may have already established a foothold by stealing credentials or changing settings. In September 2021,

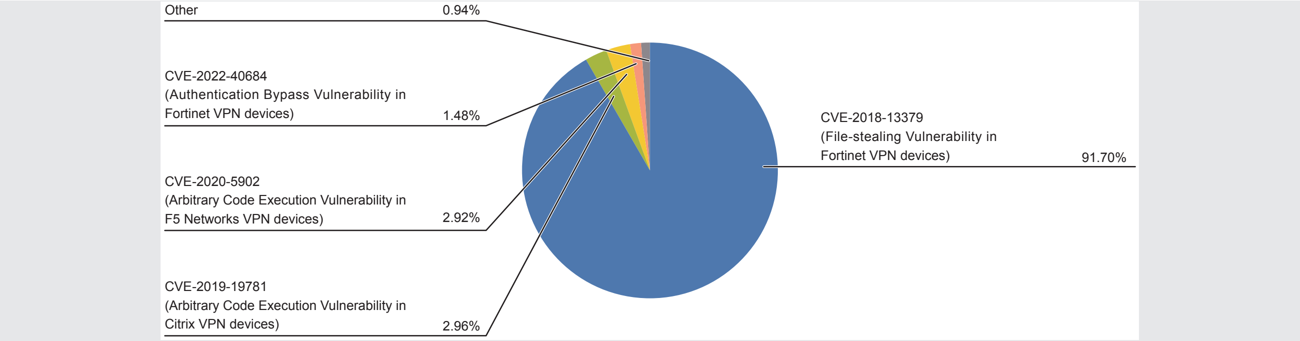

Figure 4: Detections by CVE ID (January–December 2022)

[*11] Tsurugi Municipal Handa Hospital, Tokushima Prefecture, "Tsurugi Municipal Handa Hospital Computer Virus Infection Incident: Expert Committee's Investigation Report" (https://www.handa-hospital.jp/topics/2022/0616/report_01.pdf, in Japanese).

[*12] Ministry of Health, Labor and Welfare, "13th Health/Medical/Nursing Care Information Utilization Investigative Committee, Materials of the Working Group on the Utilization of Medical Information" (https://www.mhlw.go.jp/stf/newpage_29667.html, in Japanese).

Fortinet announced that someone had disclosed credentials for some 87,000 Fortinet VPN devices worldwide[*13]. It said that the information disclosed was collected from VPN devices around the world via attacks exploiting the above-mentioned CVE-2018-13379. It has also been revealed[*12] that the information disclosed included credentials of ransomware victims in the incident calendar (Tables 1 and 2). When updating an OS or firmware, please also check for traces of unauthorized access to the VPN device and changes to device settings, and be sure to change credentials.

1.3.3 Vulnerabilities in 2022

As shown in the incident calendar (Tables 1 and 2), multiple software vulnerabilities were also published in 2022, and attacks exploiting those vulnerabilities occurred. This section covers vulnerabilities published in 2022 that our SOC observed as being exploited. Figure 5 shows a breakdown of the observations. All of the most commonly observed vulnerabilities carried the potential for remote code execution. Remote code execution (RCE) is when an attacker feeds a script containing specially crafted strings into an application (via an HTTP request, for example), causing arbitrary code to be executed on the application server that processes the input. Attackers can exploit RCE vulnerabilities to attempt various activities, including information theft, system hijacking and tampering, and malware distribution. This type of vulnerability is therefore generally considered a serious threat.

■ F5 BIG-IP Remote Code Execution Vulnerability (CVE-2022-1388)

TA BIG-IP iControl vulnerability (CVE-2022-1388) was published on May 4, 2022[*14]. BIG-IP is a family of communications control equipment made by F5 Networks. These products are deployed on corporate networks around the world. iControl is a REST API for operating BIG-IP products.

Attackers can use this vulnerability to pass specially crafted HTTP requests to the targeted BIG-IP system's iControl in order to bypass authentication and thereby gain the ability to execute arbitrary system commands with root privileges. If iControl on a BIG-IP system with this vulnerability is exposed on the Internet, attackers may be able to change device settings, which could lead

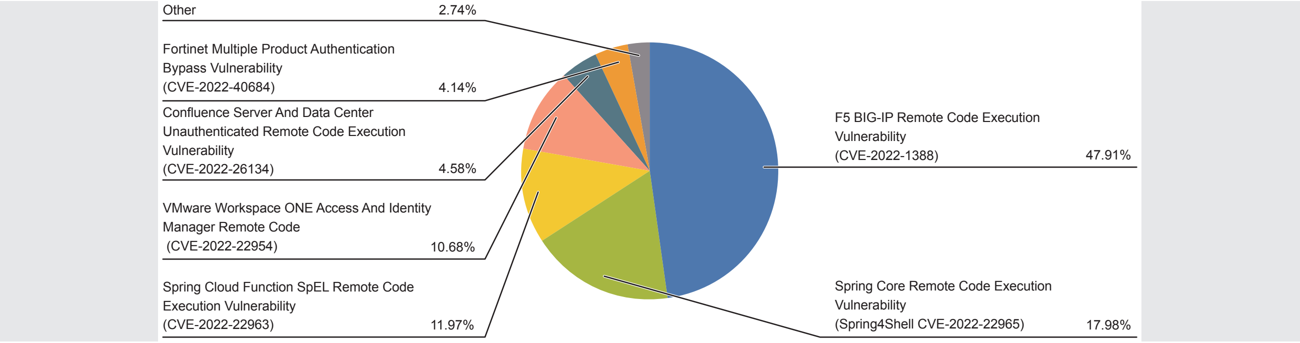

Figure 5: Breakdown of 2022 Vulnerability Exploits Observed

[*13] Fortinet, "Malicious Actor Discloses FortiGate SSL-VPN Credentials" (https://www.fortinet.com/blog/psirt-blogs/malicious-actor-discloses-fortigate-ssl-vpn-credentials).

[*14] F5 Networks, "Final - K23605346: BIG-IP iControl REST vulnerability CVE-2022-1388" (https://support.f5.com/csp/article/K23605346).

to serious damage including internal network intrusion and traffic eavesdropping.

As Figure 5 shows, this was the most commonly observed vulnerability among those published in 2022, accounting for roughly half (47.91%) of the total. Many tools for exploiting this vulnerability have been released on source code repositories such as GitHub[*15], making them available for anyone to download and run. Code for using this vulnerability is also implemented in Metasploit, an open source penetration testing tool. Attackers do abuse Metasploit, and these tools make it easier to mount attacks. The availability of such attack tools and the high popularity of the product likely explain why these attacks are so prevalent. Figure 6 shows the trend in attacks targeting this vulnerability over 2022. The observations are normalized so that the total number of attacks detected during the period sums to 100% on the vertical axis. Observations started on June 1 and peaked two weeks later on June 18. The number of attacks then declined gradually but still remained high until mid-October. The attacks came from a wide range of sources spanning 36 countries across 6 continents. So we can infer that the attacks are being carried out worldwide. Even after this activity died down in November, we continued to intermittently observe a small number of attacks.

■ Spring Core Remote Code Execution Vulnerability (Spring4Shell CVE-2022-22965)

A vulnerability related to the Spring Framework (CVE-2022-22965), commonly known as Spring4Shell, was published on March 31, 2022[*16]. The Spring Framework is a Java-based web application development framework that was open sourced by VMware as one of the Spring Projects[*17].

This vulnerability (CVE-2022-22965) exists in Spring Core (the Spring Framework's core module) and could allow remote code execution in Java applications that use the Spring Framework.

The Spring Framework accounts for a large share of Java-based web application development, but because the conditions under which this vulnerability could be executed were limited, only a small number of users were impacted.

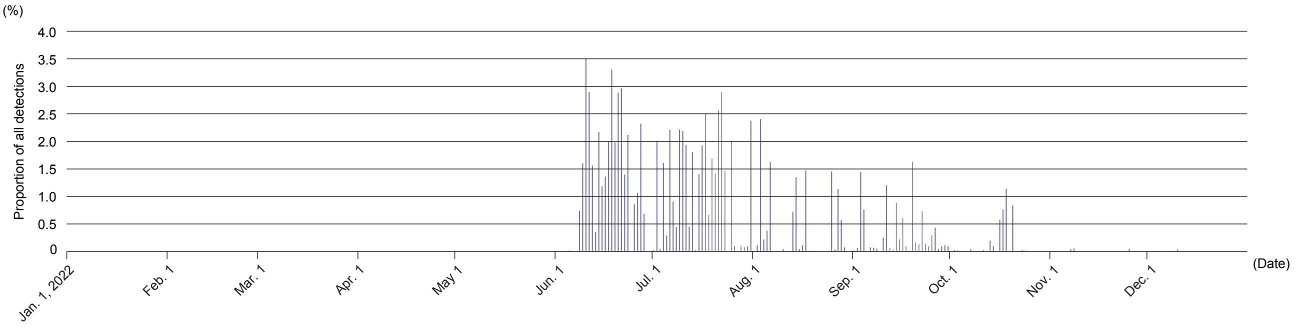

Figure 6: Observations of Attacks Targeting a BIG-IP Vulnerability (CVE-2022-1388) (January–December 2022)

[*15] GitHub (https://github.com/).

[*16] Spring, "Spring Framework RCE, Early Announcement" (https://spring.io/blog/2022/03/31/spring-framework-rce-early-announcement).

[*17] Spring (https://spring.io/projects).

When this vulnerability is used to execute remote code, HTTP requests are sent in two steps. The first HTTP request exploits the application logging functionality to write a log file to the web server containing a program (a webshell) for executing arbitrary commands. The second HTTP request then specifies the log file (webshell) created in the previous step as the URL path and sends arbitrary commands. This results in arbitrary code being executed through the webshell created in the first step. Note that Figure 7 only graphs communications used to create the webshell in the first step.

This vulnerability was the second most commonly observed among those published in 2022.

The third most commonly observed vulnerability (CVE-2022-22963) is also related to the Spring Project. It is a separate vulnerability that existed in Spring Cloud[*18]. Spring Cloud is a project concerned with the development of cloud environments, and because it has a limited user base relative to that of the core Spring Framework, the Spring Framework vulnerability (CVE-2022-22965) likely has a greater impact.

■ VMware Workspace ONE Access And Identity Manager Remote Code (CVE-2022-22954)

A vulnerability (CVE-2022-22954) related to VMWare's Workspace ONE Access, formerly known as Identity Manager, was published on April 6, 2022[*19].

Workspace ONE is a cloud-based application platform, and Workspace ONE Access is an application that manages access to workspaces.

This is a remote code execution vulnerability stemming from the way the target application's template engine processes templates. A template engine is a technology that processes input data to generate a document based on a template. Template engines are commonly used by web applications to dynamically generate HTML files. When an HTTP request containing a specially crafted character string is sent to an application server with this vulnerability and the template engine processes the input, it allows arbitrary code to be executed on the server. This sort of attack is called a server-side template injection.

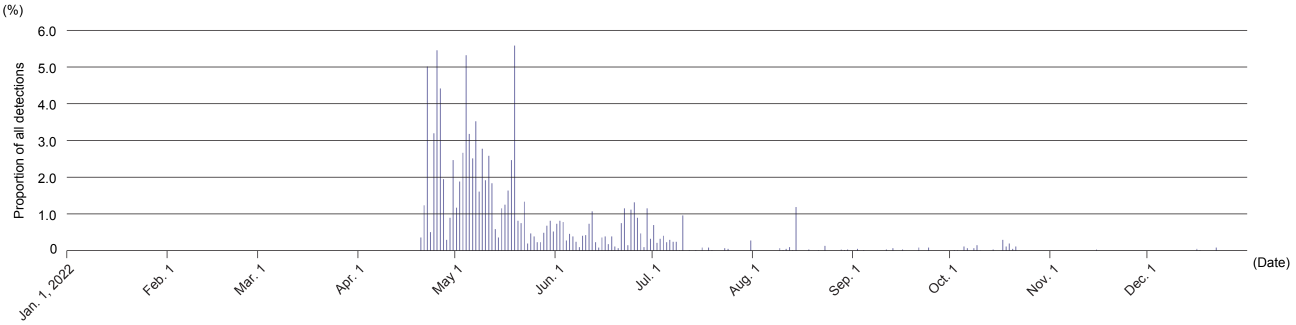

Figure 7: Observations of Attacks Targeting a Spring Framework Vulnerability (CVE-2022-22965) (January–December 2022)

[*18] VMware, "CVE-2022-22963: Remote code execution in Spring Cloud Function by malicious Spring Expression" (https://tanzu.vmware.com/security/cve-2022-22963).

[*19] VMware, "VMSA-2022-0014" (https://www.vmware.com/security/advisories/VMSA-2022-0014.html).

On May 18, the US Cybersecurity and Infrastructure Security Agency (CISA) issued Emergency Directive 22-03 instructing US government agencies to take certain actions to combat multiple vulnerabilities in VMware products, including this vulnerability[*20]. This was the only Emergency Directive issued in 2022, which highlights just how much this was seen as high-risk vulnerabilities that could pose a major threat to US government agencies.

This vulnerability was the fourth most commonly observed vulnerability among those published in 2022. As Figure 8 shows, we logged our first observation on June 5, and the attacks continued intermittently thereafter. The number of attacks spiked on August 18 to a level roughly 20 times higher than the second highest number of observations, logged on October 18. Around 99.75% of the August 18 attacks originated from a single source, while the destinations were wide ranging. This surge in attacks was also reported by Mitsui Bussan Secure Directions's SOC[*21], and can be regarded as an example of a single attacker attempting a broad, large-scale attack. In addition, there were large increases in the number of attacks every one to two months, but they originated from a different country each time. As of this writing (January 2023), we continue to observe this vulnerability being exploited.

Confluence Server And Data Center Unauthenticated Remote Code Execution Vulnerability (CVE-2022-26134)

A vulnerability related to Confluence (CVE-2022-26134) was published on June 2, 2022[*22].

Confluence is an enterprise wiki application from Atlassian, and many companies have it installed. This vulnerability is a remote code execution vulnerability in the on-premise versions of Confluence Server and Confluence Data Center. All supported versions contained the vulnerability, so a wide range of users were susceptible. Including out-of-support versions, all versions of Confluence since 1.3.0, the first version released in 2004, were vulnerable. The cloud version of Confluence Cloud is not susceptible to this vulnerability. An updated version was not available when the vulnerability was published; a fix was released the following day, June 3, 2022.

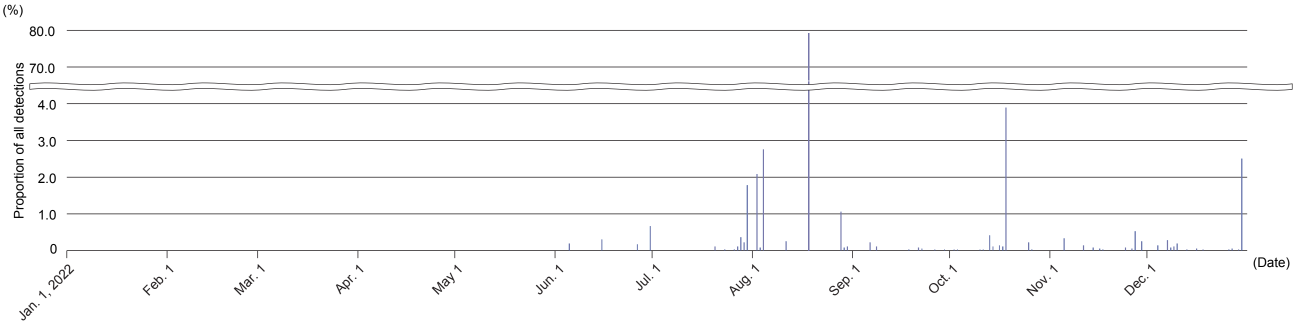

Figure 8: Observations of Attacks Targeting a Workspace ONE Access / Identity Manager Vulnerability (CVE-2022-22954) (January–December 2022)

[*20] US Cybersecurity and Infrastructure Security Agency, "Emergency Directive 22-03: Mitigate VMWare Vulnerabilities" (https://www.cisa.gov/emergency-directive-22-03).

[*21] Mistui Bussan Security Directions, "August 2022: MSBD-SOC Detection Trends and Topics" (https://www.mbsd.jp/research/20220914/20228-mbsd-soc/).

[*22] Atlassian, "Confluence Security Advisory 2022-06-02" (https://confluence.atlassian.com/doc/confluence-security-advisory-2022-06-02-1130377146.html).

The vulnerability is due to the execution of a Java-like expression language called OGNL (Object Graph Navigation Language). The execution of remote code can be triggered by injecting an OGNL expression containing a specially crafted string into an HTTP request and sending it to the targeted server. This is called OGNL injection, and this type of vulnerability is known to have caused major damage on the Apache Struts 2 web application framework in recent years.

This vulnerability was the fifth most commonly observed vulnerability among those published in 2022. As Figure 9 shows, we logged our first observation on June 18, followed by brief increases in attacks at intervals of about a month. On many of the days on which attacks increased, we observed a lot of exploit code attempting to execute the Linux id command. The id command only lists information on the user executing it, so it poses no direct threat when executed. That said, this command is often used by attackers to determine whether a target has a vulnerability, and if attackers find a vulnerability to be present, there is a risk they will carry out a malicious attack at a later time. In addition to exploit code using the id command, exploit code using the Linux wget and curl commands has also been observed. This exploit code downloads a malicious script from an external site and executes it. If this is executed on a server, it could result in malware being installed. As of this writing (January 2023), we continue to observe this vulnerability being exploited.

This section has looked at four vulnerabilities widely observed by our SOC from among those published in 2022. Susceptible product versions can be found via the URLs cited for each vulnerability[*14][*16][*19][*22]. If you are running a susceptible version of one of these applications, we recommend obtaining an updated version.

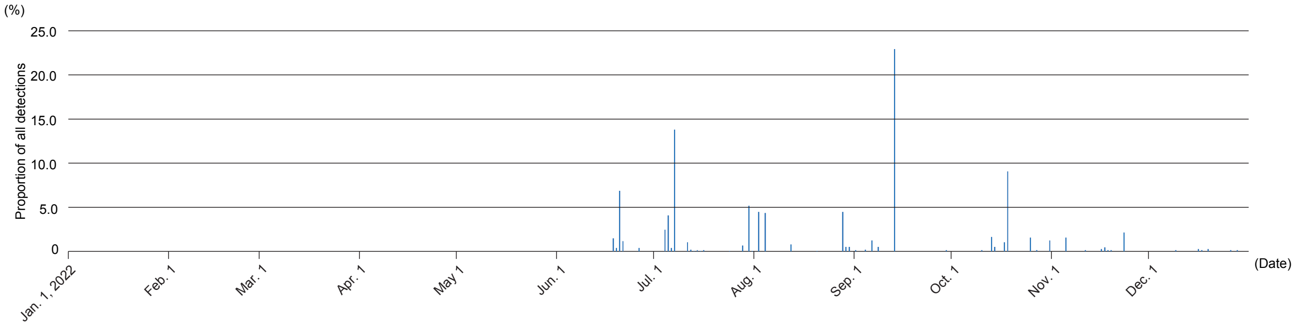

Figure 9: Observations of Attacks Targeting a Confluence Vulnerability (CVE-2022-26134) (January–December 2022)

1.4 Conclusion

This report covered security incidents that drew attention in 2022 and discussed our observations on those that our SOC analysts were focused on during the year. Our observations show that even some older attacks continue to persist. In Section 1.3.1, for example, we discussed how Emotet has seen repeated bouts of proliferation and quiescence while being updated over time, and in Section 1.3.2 we explained that a FortiOS vulnerability (CVE-2018-13379) published in 2019 was the most common one we observed among attacks targeting VPN devices. We need to remain vigilant always, not just when security topics erupt. As Section 1.3.3 discussed, meanwhile, we also observed new attacks exploiting vulnerabilities published in 2022. Constantly gathering information on vulnerabilities and updates for the products and associated services you use is crucial.

IIJ’s SOC will continue to use wizSafe Security Signal and other avenues to publish information on threats observed via our Data Analytics Platform, key security topics, and the like in the hopes that it will prove useful to you in your security responses and operations.

Eisei Honbu
Data Analysis Section, Security Operations Department, Advanced Security Division, IIJ

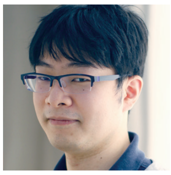

Shimpei Miyaoka
Data Analysis Section, Security Operations Department, Advanced Security Division, IIJ

Katsuhiro Tomiyama
Data Analysis Section, Security Operations Department, Advanced Security Division, IIJ

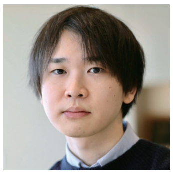

Shota Saito
Data Analysis Section, Security Operations Department, Advanced Security Division, IIJ

Interplay Between Data Centers and Electricity Markets

2.1 Introduction: Power Markets and Data Centers

IIJ's core business is in the telecommunications market, a market that has transformed considerably since telecommunications were liberalized in Japan in 1985, unleashing the principles of competition. The services themselves have also developed in many ways, with the shift from telephones to the Internet, and from fixed lines to mobiles, and market size has roughly tripled, going from 5 trillion yen in 1985 to 15 trillion yen in 2020. Over that time, liberalization has brought various benefits to users, such as cheaper service fees and the availability of a whole range of services on the Internet.

In its electricity markets, meanwhile, Japan embarked on the full liberalization of retail electricity in 2016 as part of electricity system reforms, but in comparison with the telecommunications market, it is taking longer for sufficient user benefits to appear. Japan's telecommunications market was able to follow the US model as a forerunner, and the evolution of technology has had a direct impact on communications costs (optical fiber has led to explosive growth in the amount of data that can be transported) and service content (with chip transistor density doubling every 18 months, CPU processing power has also increased explosively). Yet there was no successful national-scale model to follow for the electricity market, and from the outside looking in, it would seem that liberalization is still a trial-and-error affair here. Power generation costs depend on fossil fuel prices and investments in renewable energy power plants, and changing the market structure through technology takes time, which may be the reason why the progress of liberalization in the electricity market has been slower than in the communications market. Global events have also had a major impact, including a sharp run-up in the cost of fossil fuel needed for generating power triggered by the Ukraine crisis, and this has made it even more difficult to tell what lies ahead for the electricity market.

IIJ uses electricity to provide services, and in data centers in particular, our customers use electric power as part of colocation services. So as a substantial power consumer, we have been working with stakeholders including electric utilities and equipment vendors to ensure we can receive a stable supply of electricity, reduce costs, and move toward carbon neutrality by reducing power consumption and making use of renewable energy. No industry can sustain commercial activity without electricity, and data centers consume a lot of power, so we believe we need to look beyond our position as a single power consumer and tackle the various issues around electric power head on. Here, we discuss challenges in the electricity market from a power consumer's perspective and explain how we plan to address them.

2.2 Electricity Market Challenge 1: Electricity Costs

Of the three electricity market sectors—power generation, transmission and distribution, and retail—the 1995 revision of the Electricity Business Act, in principle, liberalized entry into the power generation sector. The liberalization of the retail sector, meanwhile, proceeded in stages, with the extra-high-voltage category being liberalized first, followed by the high-voltage category, which serves small and medium-sized factories, and then the low-voltage category, which serves residential needs, thus completing the full liberalization process. Many retail electric power companies took this opportunity to enter the market, and just as electricity prices were lowered and the effects of liberalization began to appear, the rise in electricity procurement costs driven by the recent surge in fuel prices led to the collapse of small-scale retailers, resulting in increases in electricity prices.

Electricity costs are generally said to account for 30–40% of data center operating costs (Figure 1). Comparing March 2021 and January 2023, the fuel cost-adjusted price increased by 15.82 yen/kWh (TEPCO extra-high voltage), which has meant 40–50% increase in data center costs overall. We will continue to do what we can on our end, including energy saving initiatives, but the situation is such that we cannot avoid passing some of the increases through to customers.

2.3 Electricity Market Challenge 2: Carbon Neutrality(Energy Savings and Renewable Energy)

A goal of the Paris Agreement, an international treaty on climate change, is to achieve carbon neutrality by 2050 by reducing greenhouse gas (CO2, methane, N2O, CFCs) emissions. Over 120 countries/territories have ratified the agreement, and the Japanese government also made its own "carbon neutral by 2050" declaration in October 2020. In December 2022, the government finalized its Basic Policy for the Realization of GX (Green Transformation), which clarifies its approach to and the process for, among other matters, promoting rigorous energy-saving initiatives, making renewable energy the main power source, using nuclear power, encouraging the use of hydrogen and ammonia by industry, the development of the electricity and gas markets, resource diplomacy, and the storage battery industry. Japan's Sixth Strategic Energy Plan, released in October 2021, also sets numerical targets for the 2030 energy mix, as shown in Figure 2, calling for energy savings equivalent to around 9% of the mix and a renewable energy ratio of 36–38% to total, up from 22–24% previously. As a substantial power consumer, IIJ also believes that it needs to take the initiative on energy savings and the use of renewable energy in an effort to achieve carbon neutrality.

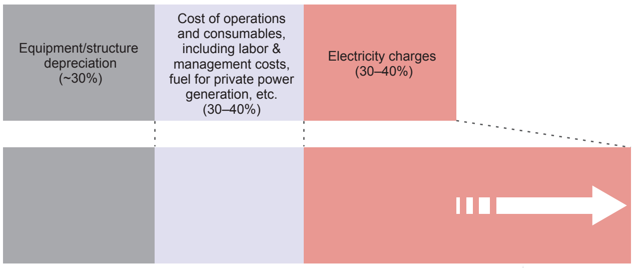
Equipment/structure depreciation (~30%)
Cost of operations and consumables, including labor & management costs, fuel for private power generation, etc. (30–40%)
Electricity charges (30–40%)

Fuel cost-adjusted price up 15.82 yen in Mar. 2023 vs. Mar. 2021 (TEPCO extra-high voltage). Total data center costs up 40–50%.

Figure 1: Typical Data Center Cost Structure[*1]

Carbon Neutral Data Center Model

To achieve carbon neutrality, we need to create a new model that organically links the generation equipment that supplies the power and the data centers that consume that power. IIJ has formulated a carbon-neutral data center reference model (Figure 3) that combines multiple power plant complexes, power storage equipment, supply and

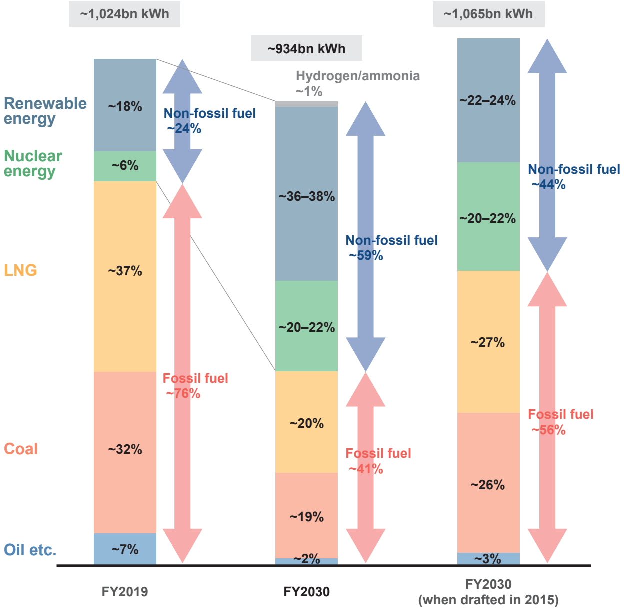

Figure 2: Energy Mix in the Sixth Strategic Energy Plan[*2]

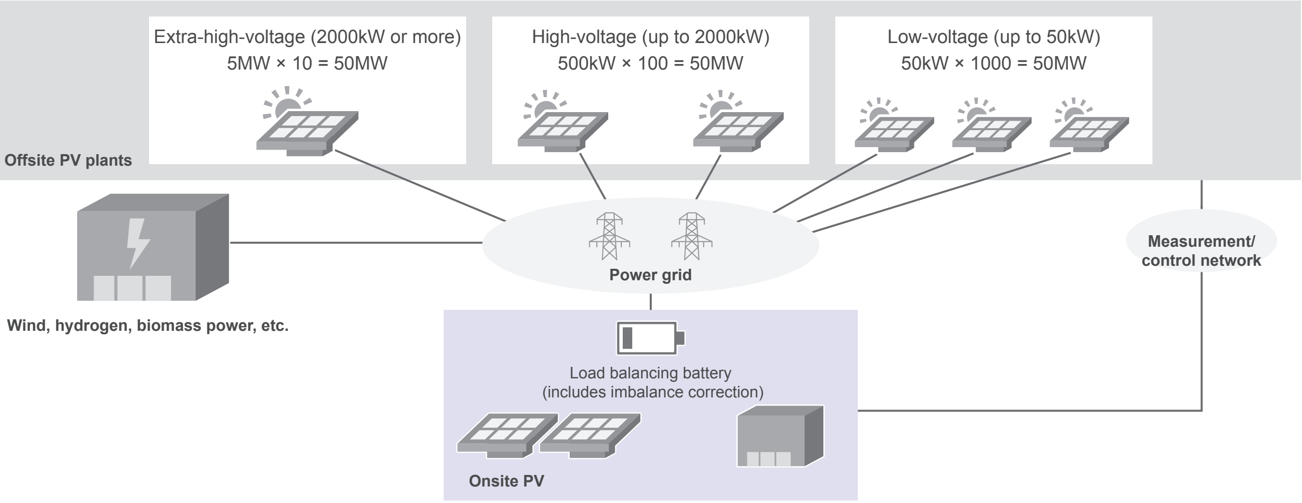
Extra-high-voltage (2000kW or more)
5MW × 10 = 50MW
High-voltage (up to 2000kW)
500kW × 100 = 50MW
Low-voltage (up to 50kW)
50kW × 1000 = 50MW
Offsite PV plants
Power grid
Measurement/ control network
Wind, hydrogen, biomass power, etc.
Load balancing battery (includes imbalance correction)
Onsite PV

Figure 3: Carbon Neutral Data Center Reference Model

[*1] Source: Adapted by IIJ from the website of Nikkei Xtech (https://xtech.nikkei.com/atcl/nxt/column/18/02096/063000007/).

[*2] Source: Adapted by IIJ from the website of the Agency for Natural Resources and Energy (https://www.enecho.meti.go.jp/about/special/johoteikyo/energykihonkeikaku_2022.html.html).

demand control, and the like, and going forward we will be conducting technical tests and working with external partners on both the business and technological fronts, and applying this model when modifying our own data centers or building new ones.

Energy Saving Trends and IIJ's Track Record

Japan's Energy Saving Act was revised in 2022 to add data centers to the list of industries subject to its benchmark system. PUE (power usage effectiveness) was adopted as the benchmark indicator, with a target of 1.4 or lower being set. Data center operators are due to report for the first time in July 2023, with the law now calling on these companies to pursue further energy savings and use electricity efficiently.

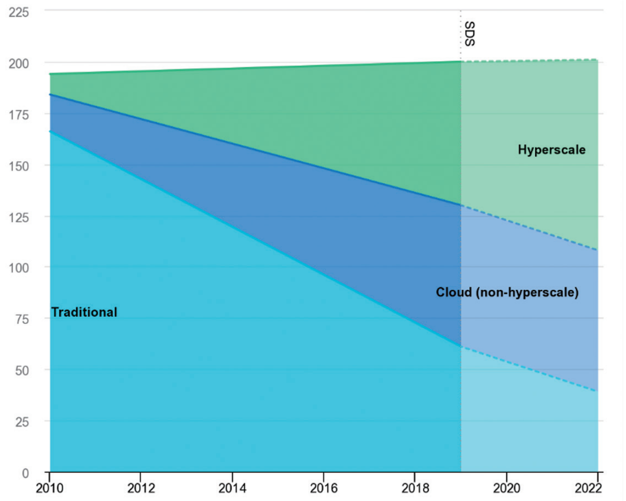

Figure 4: Global Data Center Energy Demand[*3]

PUE is found by dividing total data center facility energy consumption by IT equipment energy consumption, with a PUE closer to 1.0 indicating better efficiency. The average in Japan is said to be around 1.7. An Uptime Institute survey of data centers worldwide showed the 2022 average to be 1.55, a substantial improvement from 2.5 in 2007.

The increasing use of power by data centers has been perceived as a problem with serious implications for the environment, with people saying that data centers globally will consume 51% of the world's power by 2030. A survey jointly conducted by the Lawrence Berkeley National Laboratory in the US and others in 2020, however, showed that while data center compute capacity increased sixfold from 2010 to 2018, data center power consumption only rose 6%, from around 194TWh in 2010, or around 1% of the world's energy consumption, to 205TWh in 2018.

Figure 4 plots changes in power consumption for three types of data centers: traditional (traditional colocation), cloud

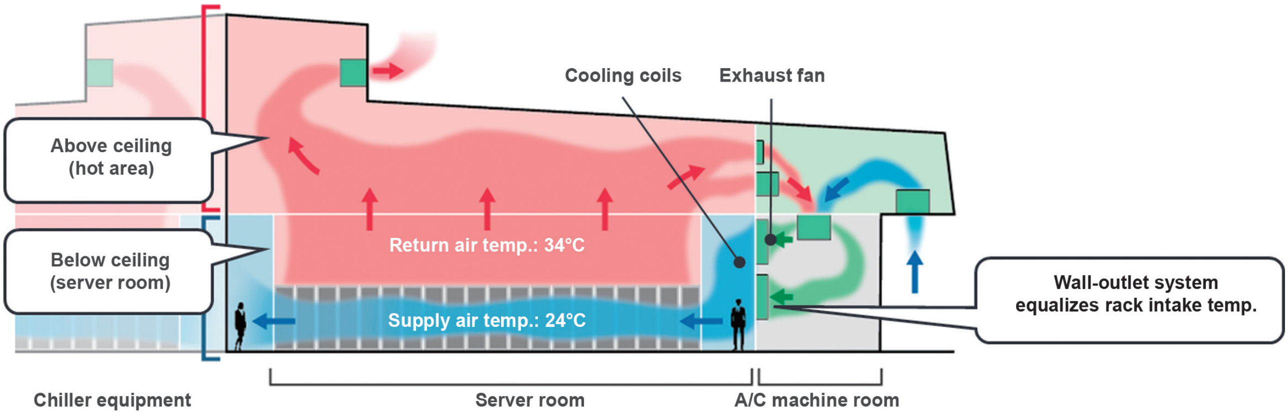

Figure 5: Outside-air Cooling Method at Shiroi Data Center Campus

*3 Source: IEA (https://www.iea.org/data-and-statistics/charts/global-data-centre-energy-demand-by-data-centre-type-2010-2022).

(non-hyperscale), and hyperscale (large-scale data centers for cloud providers). The share of power consumed by hyperscale data centers increased by nearly 30% from 2010 to 2018, and this spread of hyperscale data centers capable of processing lots of data without using much power has likely been a factor in limiting the growth in overall data center power consumption, leading to better PUE readings.

■ IIJ's Initiatives

IIJ operates its own data centers in Matsue-shi, Shimane prefecture, and Shiroi-shi, Chiba prefecture, where it has introduced energy-saving technologies to make the facilities run efficiently. Air-conditioning systems are the next biggest consumer of power in the data center behind IT equipment, and IIJ uses outside-air cooling systems (Figure 5). Standardizing on 230V power input for the IT equipment made it possible to adopt a three-phase four-wire power distribution system that can supply the 400V output of the UPS (uninterruptible power supply) to the servers without having to step it down and incur transformer losses (FIgure 6). These efforts have resulted in a PUE of 1.2.

Figures 7 and 8 show actual PUE readings. Matsue Data Center Park (Matsue DCP) opened in 2011 and has consistently achieved a PUE in the 1.2–1.3 range since 2017. Shiroi Data Center Campus (Shiroi DCC) went live in 2019, and its PUE has improved as its utilization rate has increased, and we expect it to move into the 1.3–1.4 range in FY2022.

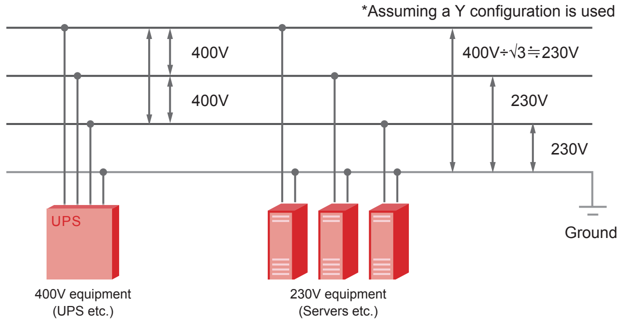

Figure 6: Three-phase Four-wire System

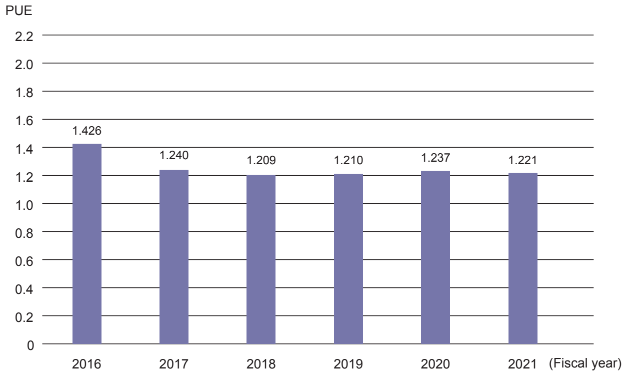

Figure 7: Matsue Data Center Park PUE (annual averages)

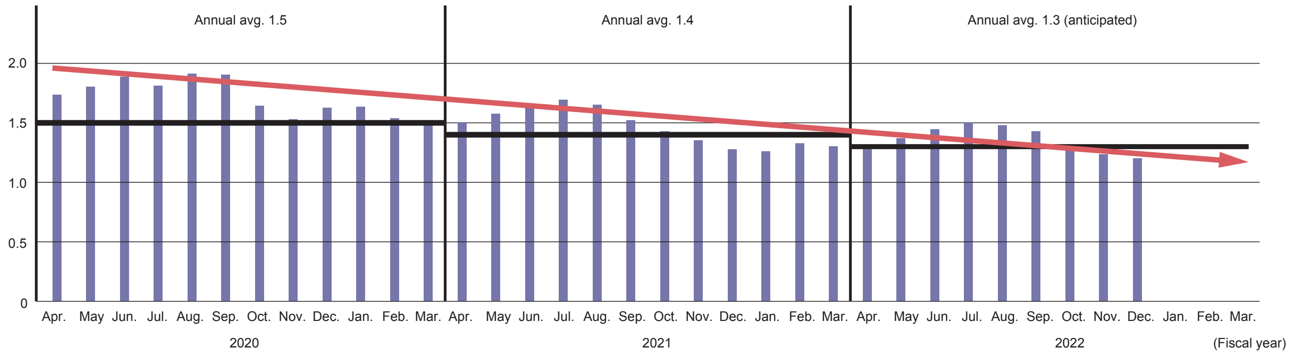

Figure 8: Shiroi Data Center Campus PUE (monthly averages)

Renewable Energy Trends and IIJ's Initiatives

In the sci-fi novel Project Hail Mary, currently being adapted into a Hollywood film, a quarter of the Sahara is covered with panels to collect huge amounts of solar energy, which, as a knock-on effect, induces natural disasters, such as frequent tornadoes in Spain, causing great harm to humanity. While this is fiction, the installation of solar panels can have adverse effects in real life too, including environmental destruction leading to landslides and the problem of how to dispose of panels after removal, so there are many issues to overcome if their use is to expand in Japan.

Renewable energy generation capacity is set to accelerate globally, however. As illustrated in Figure 9, the EIA reported in December 2022 that over the next five years through 2027, global renewable energy capacity is set to increase by 2,400GW (2.4TW), which is equivalent to the entire installed power capacity of China today.

Of all the world's industries, the IT industry has made the most progress adopting renewable energy. Figure 10 shows the top 10 procurers of renewables through PPAs (power purchase agreements), under which companies purchase renewable energy directly from power producers. Five of the 10 are data center operators (Google, Facebook, Amazon, Microsoft, QTS). The use of renewable energy can of course help with investor relations, but globally, the cost of generating renewable energy from wind and solar is coming down more than the cost of generating power conventionally using fossil fuels, so for data center operators that constantly consume large amounts of power, the shift toward renewables also reflects economic rationality.

IIJ is also pursuing renewable energy initiatives. In February 2022, we began using electricity derived from deemed renewable energy sources, which add energy attribute certificates to electricity from power utilities, at Matsue DCP. RE100 is an international initiative under which companies aim to procure 100% of the energy they consume in their business activities from renewables. In October 2022, the technical criteria for RE100 members were revised such that with respect to power purchased from retailers and energy attribute certificates, only power procured from generation facilities within 15 years of the facility being commissioned or expanded/upgraded is recognized as renewable energy for RE100 purposes, and so looking ahead, companies are likely to increasingly install new onsite generation facilities to meet their own power needs and introduce offsite corporate PPAs. Moreover, Japan's Energy Saving Act is slated for a revision in April 2023, adding new measures on transitioning to non-fossil fuel energy and requiring

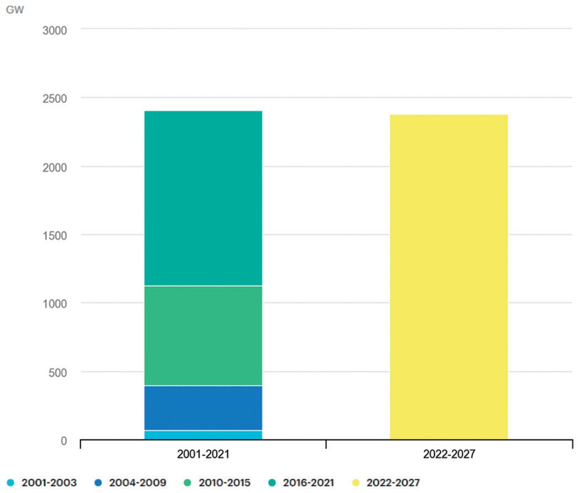

Figure 9: Use of Renewable Energy to Accelerate over 2021–2027[*4]

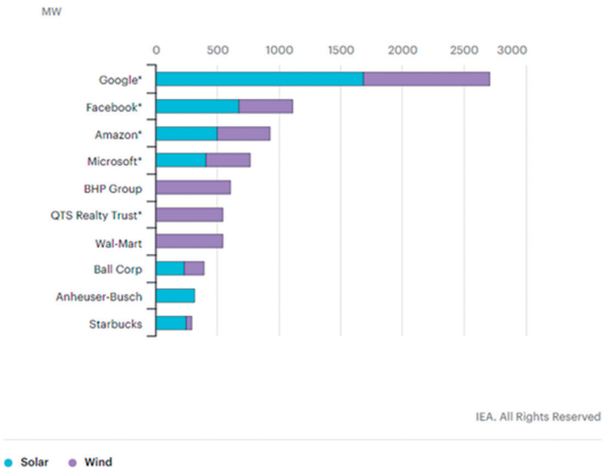

Figure 10: Top 10 PPA Procurers (2019)[*5]

*4 Source: IEA Renewables 2022 (https://www.iea.org/reports/renewables-2022/executive-summary).

*5 Source: IEA (https://www.iea.org/commentaries/data-centres-and-energy-from-global-headlines-to-local-headaches).

companies to report on their usage of non-fossil fuel energy, so an increasing number of companies can be expected to put effort into adopting renewables.

Onsite solar power generation systems, depicted in Figure 11, are scheduled to go into operation at both Matsue DCP and Shiroi DCC during fiscal 2022, but one issue is that only a small amount of power (a few percent) relative to total data center usage can be generated from onsite systems. With the cost of generating renewable energy dropping every year, our next step will be to look at engaging in offsite corporate PPAs and ownership of power generation plants that deliver power through the grid.

2.4 Electricity Market Challenge 3: Creating New Markets for the Stable Supply of Electricity

■ Entry into New Markets, the Capacity Market

The Great East Japan Earthquake of 2011 marked a turning point, prompting reforms to the electric power system that have been underway in stages since 2015 with the objectives of ensuring a stable supply of electricity, curbing electricity charges, and expanding options for power consumers and business opportunities for operators. New markets have been created as shown in Figure 12.

The capacity market is one of those new markets. Electricity market liberalization and the decline in non-renewable power

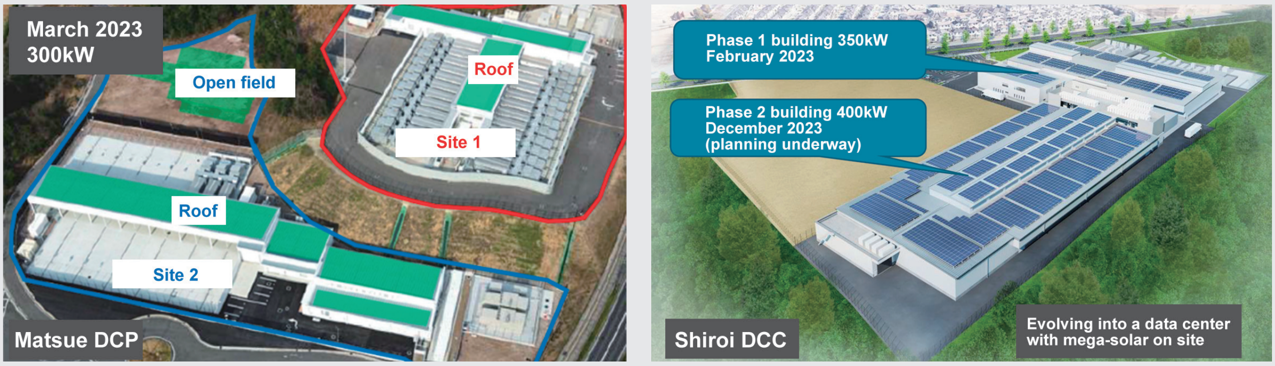

Figure 11: Installing Onsite Solar Power Facilities

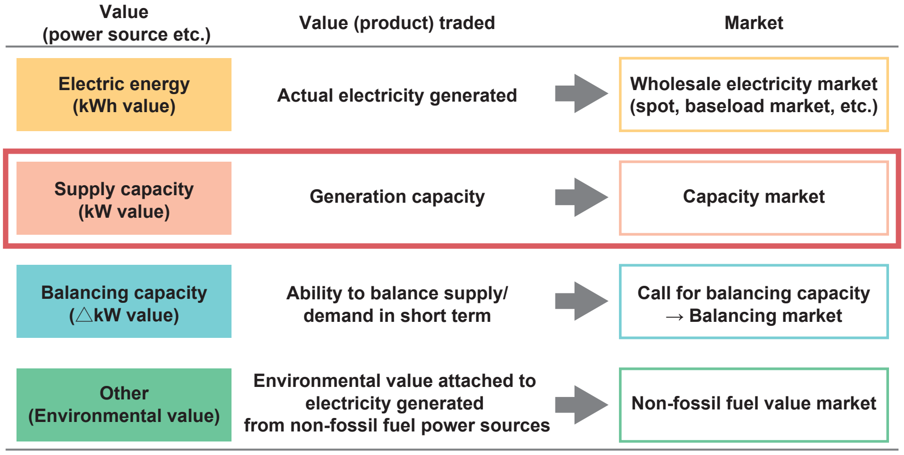

Value (power source etc.)	Value (product) traded		Market
Electric energy (kWh value)	Actual electricity generated	→	Wholesale electricity market (spot, baseload market, etc.)
Supply capacity (kW value)	Generation capacity	→	Capacity market
Balancing capacity (△kW value)	Ability to balance supply/ demand in short term	→	Call for balancing capacity → Balancing market
Other (Environmental value)	Environmental value attached to electricity generated from non-fossil fuel power sources	→	Non-fossil fuel value market

Figure 12: New Trading Markets[*6]

[*6] Source: Adapted by IIJ from the website of the Organization for Cross-regional Coordination of Transmission Operators, Japan (https://www.occto.or.jp/capacity-market/shikumi/capacity-market).

plant utilization and market prices owing to the growing use of renewable electricity bring about the risk of investments in new generation capacity not moving forward because the prospects of recouping those investments is less predictable. To reduce this risk and ensure electricity supply capacity is available into the future, the capacity market allows participants to trade in supply capacity (kW) instead of amount generated (kWh).

Participants in the capacity market trade in stable power supplies, variable power supplies, and demand-driven power supplies. Within the demand-driven category, from fiscal 2024 IIJ will be supplying electricity as part of a virtual power plant (VPP)[*7] aggregated by Kansai Electric Power Co. At Shiroi DCC, we will use the lithium-ion storage batteries, installed to smooth out summertime air-conditioner power consumption, to facilitate our demand response (DR)[*8], one means of controlling the electricity demand/supply balance for a VPP. We will use the batteries' surplus capacity as well as onsite solar power generation to respond to requests to reduce grid power consumption, thus receiving payments from the aggregator and reducing our data center operating costs (Figure 13).

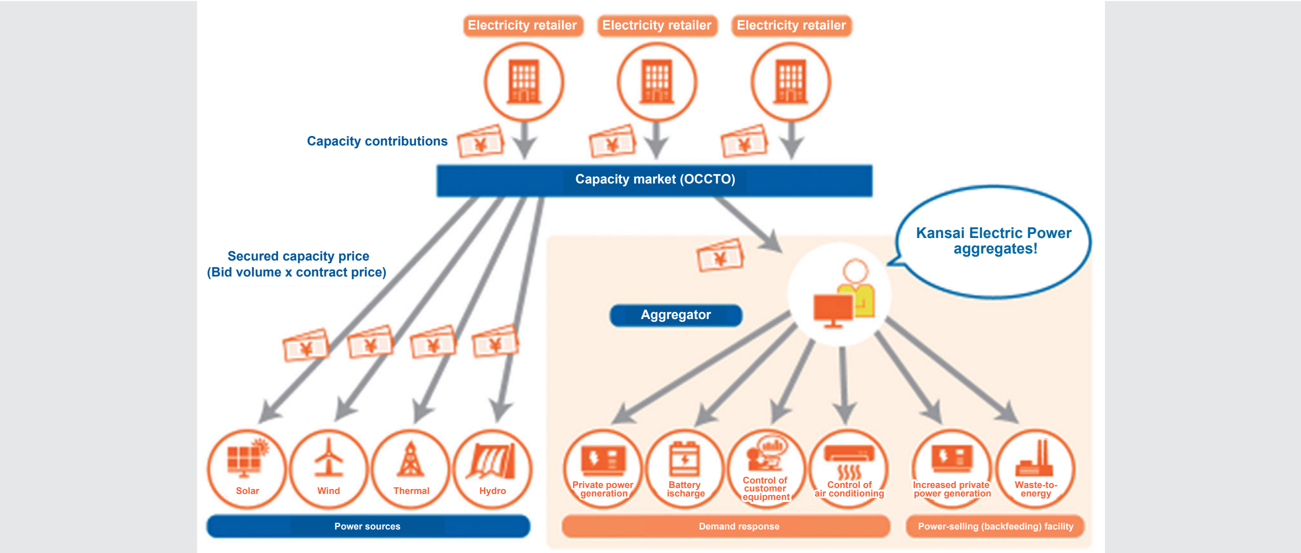

Figure 13: Participating in the Capacity Market[*9]

[*7] VPP (virtual power plant): A collection of regionally dispersed power generation facilities—storage batteries owned by companies and local governments, small power generation facilities, and the like—controlled by an aggregator in an integrated manner so as to function like a single integrated power plant.

[*8] DR (demand response): DR refers to controlling the electricity supply/demand balance by controlling the electricity usage of consumers through methods such as setting different electricity rates for different times of the day and paying consumers who refrain from using electricity during peak hours.

[*9] Adapted by IIJ from the website of Kansai Electric Power (https://www.kepco.co.jp/energy_supply/energy/vpp/market.html).

2.5 Conclusion

This article has discussed how we are addressing issues in the electricity market as a consumer of power in our data centers. In closing, I would like to explain what IIJ hopes to provide to data center users on the power supply front.

First is the ability to visualize power consumption. Going forward, data center users will likely be called on to save more energy than ever before. So users also need to know how much electricity they are using. We plan to move forward quickly with the development of a system that will provide detailed data on each customer’s current power usage.

Secondly, we hope to provide renewable energy value. Efforts to achieve carbon neutrality will continue to advance, and in that context, we are looking at the prospect of building a platform that will enable us to provide the environmental value of renewables to our customers based on the ability to visualize data center users’ renewable energy breakdown by type and as a percentage of total energy consumed.

As an operator of data centers that use large amounts of electricity, we will continue striving toward carbon neutrality as part of our social responsibility, and we hope to be able to continue reporting on the positive results of our efforts.

Isao Kubo

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

The New IIJ Studio TOKYO’s Bridge to the Future

3.1 Introduction

While IIJ is known as an ISP, we have actually been in the video delivery business since the 1990s. Over the past few years, we have been providing large-scale content distribution services like the IIJ MediaSphere Service, and we have been expanding our services to meet the increasingly diverse needs of customers. Online video delivery has increased since 2020 brought major changes in the social landscape, and many companies now deliver video in line with that trend. IIJ also streams its own financial results presentations and other events to an external audience, and we have received high praise for video and audio quality, which has an impact on corporate branding, and customers have told us they would like to achieve the same. When streaming earnings announcements, we make full use of the knowledge and experience of onsite staff to achieve a stable, high-quality stream. We have cameras and switchers (devices for switching video signals) installed temporarily in the conference room, and these are operated by members of our PR team and multiple other teams. But because the events were recorded in an ordinary conference room, we often had to deal with noise from outside, sudden accidents, and the like. IIJ also streams Japan’s largest classical music festival, the Spring Festival in Tokyo, every spring, and has offered a paid live stream of the event since 2021. In 2021, we set up a temporary streaming center at the Tokyo Bunka Kaikan, a concert hall in Ueno, which received video from multiple venues and streamed it out. Quite apart from the streaming, this setup involved a lot of effort in terms of transporting equipment to and building the streaming center. So in 2022, we built a streaming center in the IIJ Iidabashi office, from where we controlled remote IP cameras located in the Tokyo Bunka Kaikan hall. Under this remote production streaming setup, video from the venues was transmitted to the Iidabashi sub-control room and streamed from there.

Even so, we had to go through the process of building a streaming center in both 2021 and 2022, and various tasks had to be dealt with, like finding a location and procuring equipment. In view of growing demand (from both inside and outside of IIJ) for video delivery, and to enable us to meet the needs of as many customers as possible, we started to look at the prospect of building a permanent streaming center and a permanent studio capable of stably delivering high-quality video.

And so it was that in October 2022, IIJ Studio TOKYO came to life in Iidabashi. But just having a permanent studio does not guarantee high-quality video streaming. Building a stable video production and streaming environment and getting it up and running takes a lot of time and effort. It involves long-term aging and verification testing, and it takes teamwork and operators with skill and experience.

As few of our staff at IIJ had any knowledge of video production, we had to start from scratch when it came to training operators, learning how to handle equipment and wind cables, and determining what level of video and audio quality was required.

In FY2022, we were focused on recording and delivering video internally, trying to learn about the equipment and devices and how to use them, building a system of operations, and so forth. But we would have been complacent to stop at this internal setup. We hope to turn this into a commercial service next fiscal year or beyond, and while we build our studio, we are also working to improve our studio operations by talking to people outside of the company and seeking feedback from within. To this end, we are engaged in dialogue with people both internal and external to IIJ and increasingly collaborating with external partners.

We also wanted to take on the video industry challenges discussed in this article. Thinking about what sort of studio IIJ should create and how it would be unique to IIJ, we designed it to be capable of handling remote production based on IT and IP and to be capable of live video streaming using mobile communication lines.

Below, we discuss the benefits and challenges in using IP in modern video production and take a look at the IIJ Studio TOKYO facilities.

3.2 Overview of IIJ Studio TOKYO Facilities

Here, we discuss the IIJ Studio TOKYO facilities.

As Figure 1 shows, IIJ Studio TOKYO houses a studio (Photo 1) and six rooms. The rooms are interconnected via a 10Gbps or 1Gbps IP network. The advantage of using IP for the studio system is that you can connect to the studio sub-control room (Photo 2) via a single optical fiber cable. For example, you can easily turn a conference room on

Photo 1: Studio

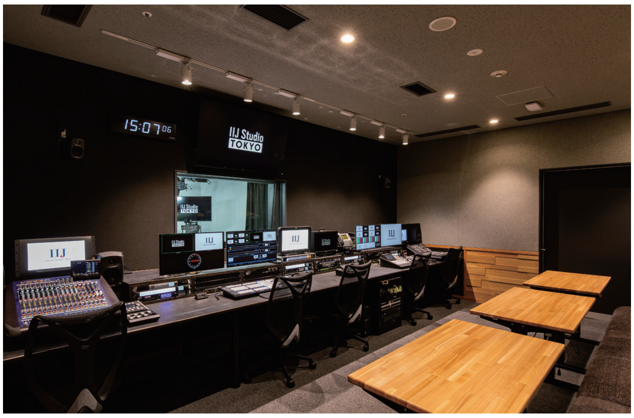

Photo 2: Studio Sub-control Room

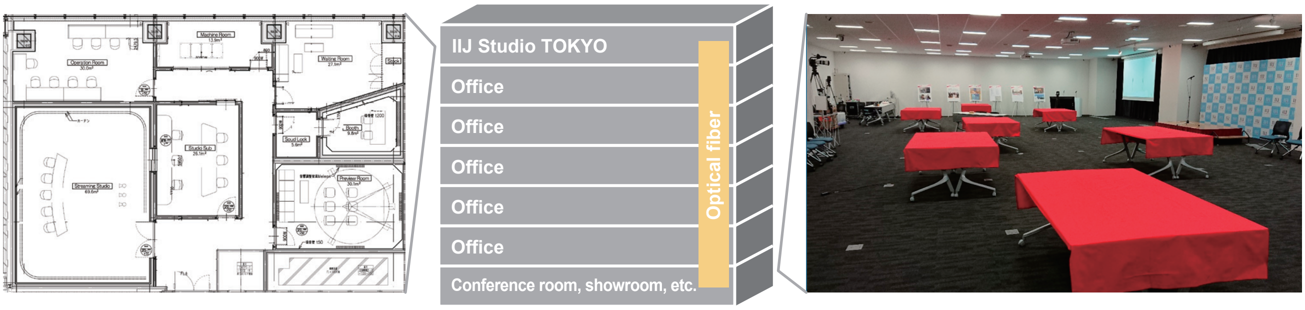

Figure 1: Studio Layout and Network Configuration

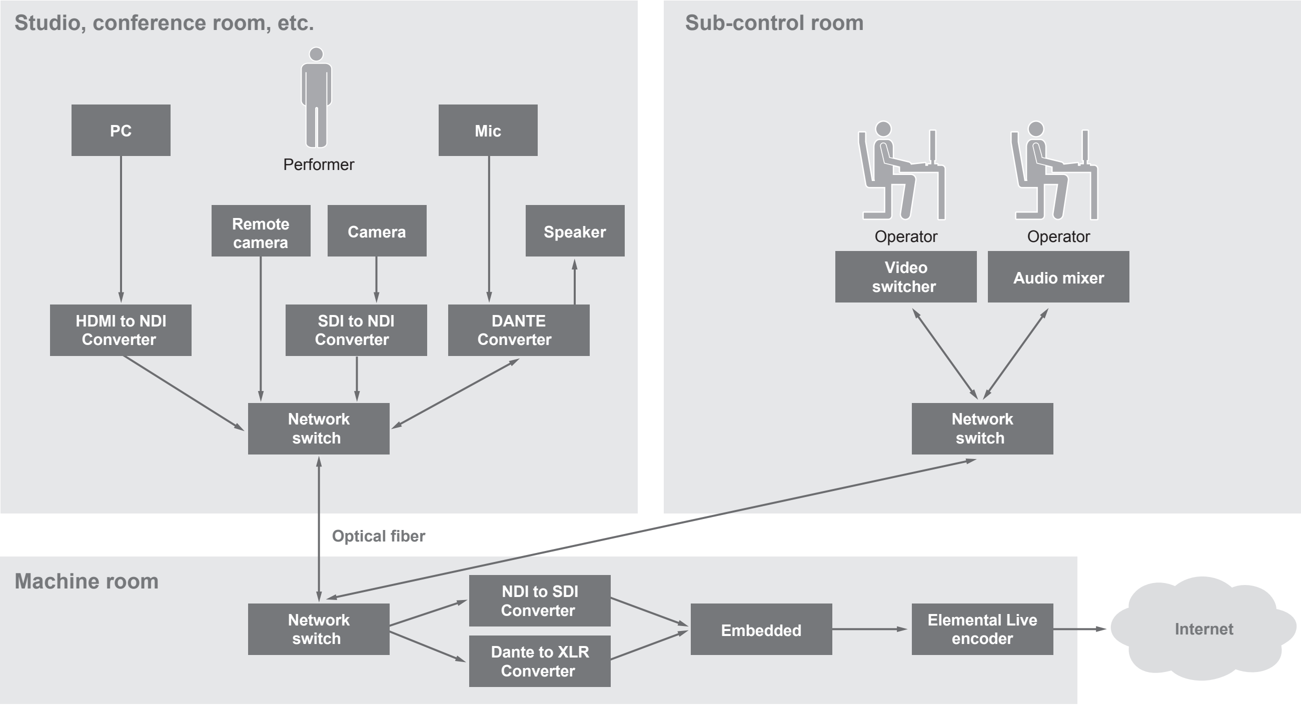

Figure 2: Example Configuration
*Figure assumes only IP is used. SDI can also be used in combination with IP.

another floor into a temporary studio by installing a camera in there. You also have the flexibility to set up a shoot and venue according to your objectives, which includes factors such as the scale of the shoot and whether an audience will be present. Adding in the elements of remote production makes it possible to centrally manage multiple camera angles and color adjustments, which require a lot of time to set up, from within the studio sub-control room, and this helps to reduce the number of onsite cameramen and operators and lighten the workload.

■ Studio

■ Cameras

A total of six cameras are used for filming in the studio (two Sony professional camcorders and four Panasonic remote cameras), so the video produced can use multiple camera angles. The basic concept is to use the remote cameras as the main sources, which helps address camera operator shortages as well. The studio also has white walls and black curtains, as well as a green screen/mat setup to facilitate chroma-key compositing, making it possible to handle a variety of situations to suit users' needs. Lighting is of course indispensable in any shoot, and this is connected via Bluetooth, allowing luminance and color temperature to be controlled from an iPad.

■ Microphones

For stability purposes and to prevent crosstalk, we use Type A radio microphones, which require a license. We use a combination of pin and handheld microphones depending on the objectives of the shoot. An omnidirectional condenser microphone is also installed on the ceiling to allow for any problems, so sound and voice can be recorded reliably when in the studio.

■ Studio Sub-control Room

■ Switchers

In the studio sub-control room, we use the NewTek TriCaster2 Elite as an IP switcher. It can take eight SDI inputs and 32 NDI inputs. So by cutting down on cabling, we have a greater degree of freedom when it comes to where we set up devices that transmit video sources and how many of them we have.

By using virtual sets, one key feature of software switchers, we can also add a sense of movement to a scene using a single camera, opening up a broader scope of video production possibilities. We also have a Panasonic switcher to use as a sub-switcher, making it possible to perform multiple operations simultaneously, with the main switcher used for the studio shoot and the sub-switcher used for remote production.

■ Digital Mixer

We use Dante for the studio's audio protocols. Using Dante devices makes it easy to change routing configurations via a PC app with Dante Controller or a digital audio mixer, so the required sound source can be routed to the required location quickly without having to fiddle with the physical cabling.

■ Machine Room

We have cut down on unnecessary cabling by consolidating all video signals in the machine room. Video can be efficiently routed to the room or device where it is needed. To make it easy to record live from external venues, we also have receivers for LiveU compact video transmission units in place, which we envision using for events like the Spring Festival. We also have an Elemental Live (encoder for Internet streaming) installed, which can be used for streaming from the studio and video input from external venues. And the communications environment takes advantage of IIJ's strengths, with a dedicated 10Gbps Internet line connected directly to the IIJ backbone.

■ Recording Booth

IIJ Studio TOKYO has a dedicated audio recording booth, which makes it possible to add narration to video being recorded in the studio or pre-record the moderator's voice for a webinar, for example. In order to provide an optimal environment for audio recording, the floor and walls are designed to reduce echo and reverberation. We have also taken care with the shape of the room and the positioning of sound-absorbing materials, and replaced the existing air conditioner with a dedicated duct, as a result of which you really notice the change in sound quality the moment you enter the room.

■ Preview Room

Productions can be previewed on a 100-inch screen with 7.1.4ch Dolby Atmos compatible sound equipment. The room can also be used as a performer greenroom or a location for shooting interviews and the like.

■ Greenroom

This is a place to relax before a performance, with a view overlooking the Kagurazaka district. The room also has a dressing table, changing curtains, a large storage space, and a refrigerator. This room, too, can be used to shoot interviews, discussions, and the like. The space is suitable for shooting still photos as well as video.

■ Operations Room

The operations room provides a work place for studio operations, such as storing and testing studio equipment, and editing recorded media. For large-scale streaming events, such as the Spring Festival, multiple streams need to be transmitted and monitored simultaneously. We have two sets of equipment in the sub-control room and two in the operations room, allowing up to four simultaneous streams. Large monitors are therefore installed in this room to display streaming status.

Photo 3: Recording Booth

Photo 5: Greenroom

Photo 4: Preview Room

Photo 6: Operations Room

3.3 Why IP? Video Industry Challenges and the Benefits of IP

The mainstream approach in the video industry has long been to produce videos using baseband signals (composite signals, SDI signals, etc.), and this baseband switcher (electronic circuit design) has been used for live video production for over 50 years without change. Back when analog circuit switchers were in use, it would take several hours after the power was turned on for the signals to stabilize and the video levels to stop changing. Decades later, the advent of digital switchers made it possible to have things operating stably as soon as the equipment was powered on. Now with increasingly high video resolutions such as 8K entering the mix, we are starting to come up against the limits of bandwidth and cable length for the 12G coaxial cables (SDI cables) that connect to the switchers.

In the IP space, meanwhile, bandwidth has increased dramatically, going from 10G to 25G and then to 100G, so a gap has opened up in terms of the pace of evolution. Live video production is thus now, after several decades, entering an era of change.

With conventional SDI, you can only transmit in one direction, either in or out, whereas with IP video transmission, you can send and receive multiple videos over a single cable. Compression technology also makes video readily transmissible, providing design flexibility with limited resources. As scale increases, so too does the impact on system building in terms of factors such as routing switchers (a device that distributes video among devices) and number of cables.

Compressed video is already used in the area of video post-production. Since 2000 or so, post-production has been transitioning from linear editing (baseband editing) to PC-based non-linear editing (file-based editing), with tools such as Avid Media Composer and Apple Final Cut Pro, and non-linear editing is now well established in this space. Yet In the area of live video production, reliability has not really been established nor has expertise been built up, so the reality is that compressed video and IP video are not a major part of live video production. The following factors explain why IP and PC tools are not widely used in live video production.

Photo 7: Coaxial Cable

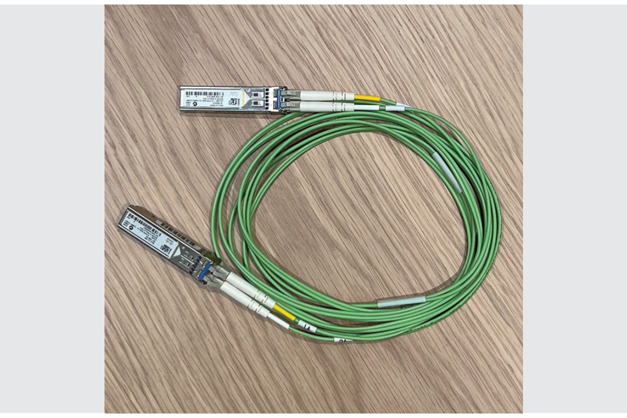

Photo 8: Optical Cable

- **Disadvantages and challenges in an IP studio, including NDI**
  - Lack of IT engineers
  - Unlike with SDI, simply connecting is not enough to enable signal transfers, systems need to be configured
  - Lack of experience when it comes to stability concerns, monitoring complexity, etc.
  - Lack of familiar methodologies or mature lineup of cabling, peripheral equipment, etc.
  - Lack of information

At IIJ, our video experts and our IP experts are working together to address the above issues with the aim of providing a stable and user-friendly environment.

A number of protocols are typically used in live video production. In our studio, we adopted NDI (Network Device Interface), which combines the best parts of compression and IP technologies. In the next section, we explain the features and advantages of NDI.

3.4 Advantages of NDI (Network Device Interface)

NDI supports 8-bit up to 12-bit signals. Its support for an alpha channel (transparency in addition to RGB) makes it flexible enough for compositing work, and it works very well with non-linear editing machines, allowing smooth and easy handover to post-production workflows. With a view to deploying it in remote production and the like, we see NDI as offering strong future potential, versatility, and cost performance. NDI also works seamlessly with a variety of devices and OSs, and it can be used on Teams apps and smartphones, making it highly versatile. The protocol allows for unobstructed communications between devices in broadcast settings all the way down to ordinary user applications, and so we believed it would be the perfect protocol for experiencing the convenience that IP offers. This is why we adopted it as the main protocol for IIJ Studio TOKYO.

	SDI	SMPTE 2022	SMPTE_2110	NDI
Compresion	×	×	○	○ NDI Codec（DCT）/NDIHX
Alpha channel	×	×	○	○
HD（1080/59.94i）Data Rate	>1.5 Gbit/s	>1.5 Gbit/s	>1.5 Gbit/s	>100 Mbit/s
UHD（2160/60p）Data Rate	>12 Gbit/s	>12 Gbit/s	>12 Gbit/s	>400 Mbit/s

Table 1: Comparison of SDI and IP Video

Even when people appear remotely via Zoom or a smartphone camera in TV productions, current practice is to convert the IP video to SDI, but it would likely be more efficient to produce programs using the unconverted IP. Producing video based mainly on IP is no easy task, however. Overhauling the baseband expertise, methodology, stability, and operations people have cultivated over many decades will require time, persistence, and the understanding of everyone involved.

In my previous job, I spent nearly a decade working to help people transition from linear editing to file-based editing and to make that change stick (demonstrations, system building, and aftermarket service), so I can see that the shift to IP in live video production is no easy task.

3.5 Video Production in the Future

Looking ahead, I think production environments will be built in the cloud, and that it will be possible to handle live video production without SDI and other specialized interfaces. But instead of everything being done in the cloud, my sense is that parts of the process that it is more beneficial to handle in the cloud will be brought in and managed centrally, and that there will be increasingly smooth integration with post-production environments as well. The NDI protocol works well in low-CPU, 1G network cloud environments and makes it possible to pass video around for internal processing in a convenient manner at low cost.

At IIJ Studio TOKYO, we first set up an IP-based production environment that is close to what a local environment offers, and we have started working on IT-driven live video production efforts that match the working environment, which also gives us an opportunity to experience stability and latency issues that can also cause bottlenecks in a cloud environment. At IIJ's Shiroi Data Center Campus, too, we have opened a research facility (Shiroi Wireless Campus) that provides an opportunity to experience the image quality and amount of latency that occurs with 4K NDI transmitted over local 5G.

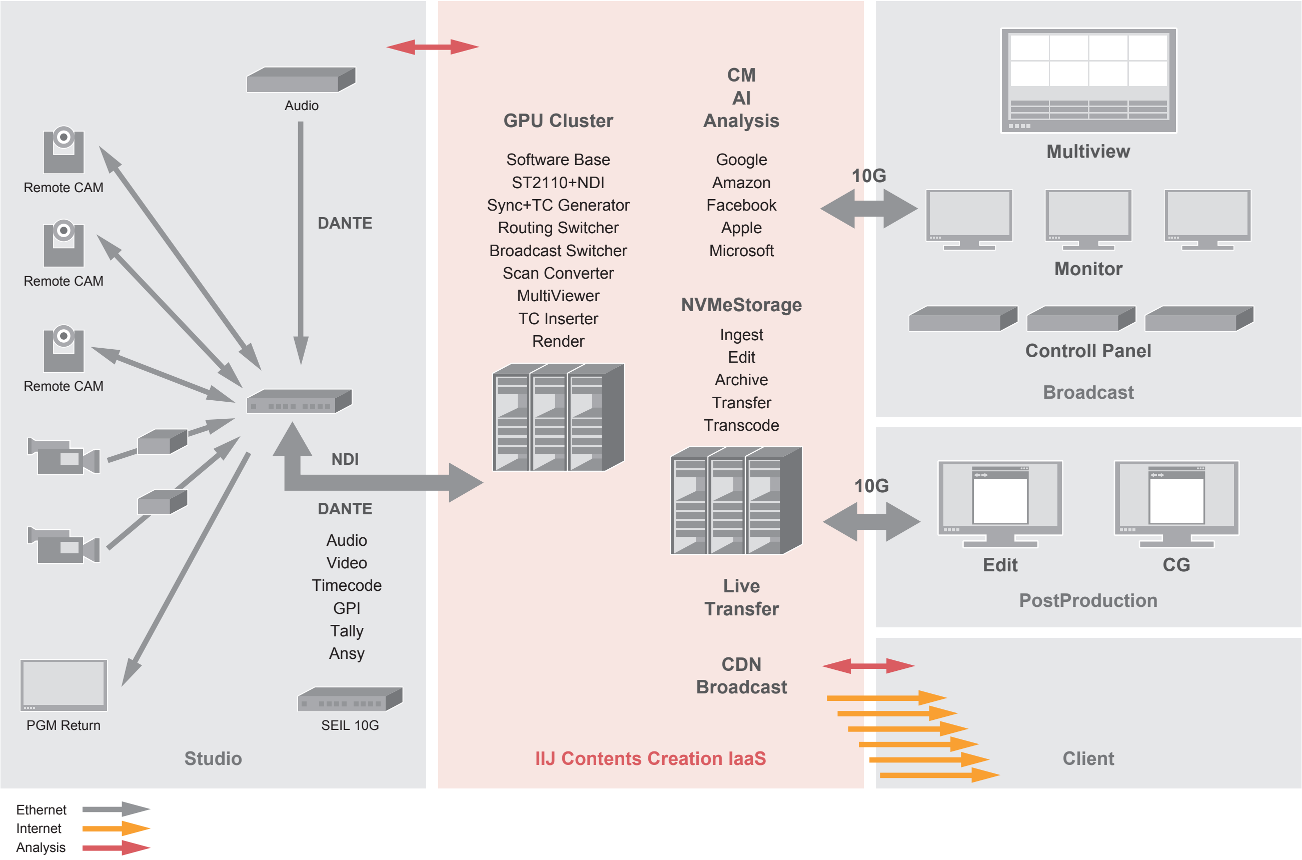

Figure 3: Conceptual Illustration of Future IT Broadcast Station

3.6 Past Accomplishments and Initiatives

2019 Accomplishments

We participated in multiple proof-of-concept projects in ST 2110 uncompressed remote production for broadcasters. Manufacturers brought their broadcasting equipment into the studio and we checked the video produced while working through the standards and monitoring communications status, and we were able to identify and share a lot of issues and information.

2020 Accomplishments

We built a 4K NDI transport demo facility that uses local 5G at Shiroi Data Center Campus, and we made the facility available so that people can experience the features listed below for themselves. This is something we recommend experiencing for people dealing with environments in which remote cameras are frequently moved, such as stadiums and factories, and for people dealing with security issues.

Photo 9: Shiroi Wireless Campus
4K NDI Image Quality and Low Latency with Local 5G

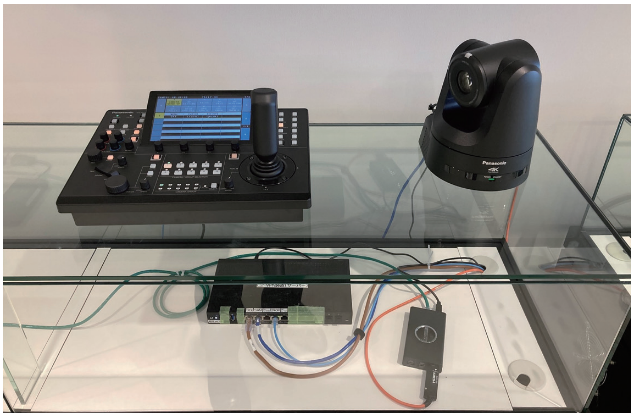

Photo 10: Panasonic PTZ Remote Controller
and Remote Camera

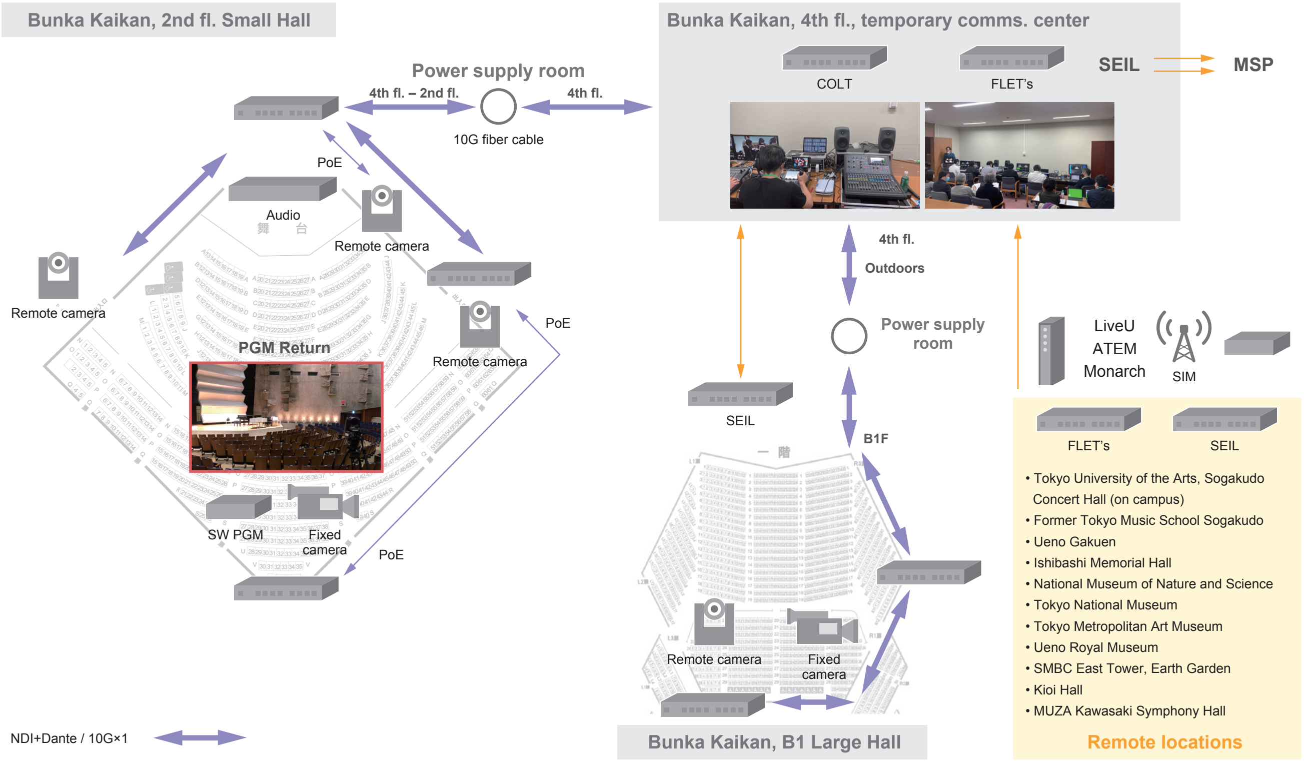
Bunka Kaikan, 2nd fl. Small Hall
Power supply room
4th fl. – 2nd fl.
4th fl.
10G fiber cable
PoE
Audio
Remote camera
Remote camera
PGM Return
Remote camera
PoE
SW PGM
Fixed camera
PoE
NDI+Dante / 10G×1
Bunka Kaikan, 4th fl., temporary comms. center
COLT
FLET's
SEIL
MSP
4th fl.
Outdoors
Power supply room
SEIL
B1F
Remote camera
Fixed camera
Bunka Kaikan, B1 Large Hall
LiveU
ATEM
Monarch
SIM
FLET's
SEIL
• Tokyo University of the Arts, Sogakudo Concert Hall (on campus)
• Former Tokyo Music School Sogakudo
• Ueno Gakuen
• Ishibashi Memorial Hall
• National Museum of Nature and Science
• Tokyo National Museum
• Tokyo Metropolitan Art Museum
• Ueno Royal Museum
• SMBC East Tower, Earth Garden
• Kioi Hall
• MUZA Kawasaki Symphony Hall
Remote locations

Figure 4: Configuration of Network at Bunka Kaikan

- **Bi-directional network offering low latency and high image quality**
- **Stable communications quality due to use of licensed frequency bands**
- **Security and lower latency compared with Wi-Fi**

🟩 2021 Accomplishments

At the 2021 Spring Festival, we built a temporary sub-control room that used NDI inside the Bunka Kaikan, one of the performance venues, and set up remote production equipment to stream the performances. The venue and the sub-control room were connected solely via a LAN cable, and we set up our infrastructure so that the video, audio, and tally ("on-air") signals from the three remote cameras in the venue could be transported via a single cable. To achieve the same cabling with ordinary baseband, you need to run a dedicated cable to each device, which makes the process more laborious and prone to connection errors, but IP makes it easy to reduce connection errors and lighten the workload. The network and camera equipment were also stable, and we determined the NDI image quality, response speed, and the like to be suitable for use in an actual production environment, so this was IIJ's first real-world NDI deployment. But because we were not covering every single venue across the entire event period, we had to lay and remove our power supply and network cabling for each performance we covered. We also had to shut down the server equipment pretty much every day, and so our first time streaming the event impressed upon us the need for a permanent sub-control room to reduce the work involved and facilitate more stable operations.

🟩 2022 Accomplishments

For the 2022 Spring Festival, we simultaneously streamed remote productions from four venues using SRT/H.265. This differs from NDI in that latency is high, so we set up two SRT connection modes in advance: a camera adjustment mode that provides good responsiveness but allows frame dropping and an event performance mode with a larger buffer size to facilitate stable transport. We switched modes right before the start of the performances. We set it up so that video from the venues was transported from the event in Ueno to the encoder in Iidabashi with a delay of five seconds, and we were able to produce a stable video stream this way.

Photo 11: PTZ Remote Camera

Photo 12: SEIL Router Used to Build L3VPN (left) and PTZ Remote Camera (right)

Photo 13: Temporary Sub-control Room

3.7 Difficulties and Solutions when Building IIJ Studio TOKYO

Our various trials and tribulations have made us recognize the need for and utility in having a permanent studio, and we thus decided to build one. Right after we put the studio into operation, we experienced equipment behaviors such as the TriCaster crashing. We also needed to learn the quirks of each product and how to do everything, including running updates, changing settings, and changing connection methods, but it is all up and running and very stable now. We also made fine-grained adjustments to settings to make efficient use of the entire network's bandwidth—we set part of the NDI network to use multicasting, for instance. For maintainability and fault tolerance purposes, we keep a full backup of the PC-based TriCaster boot image, and we regularly perform backups in the event of major system updates or system changes to enable a swift recovery if anything goes awry. The internally created TriCaster session data is also synced with the file server and constantly backed up, and we have standardized and documented our system so that if the TriCaster were ever to fail to boot at some point, we can immediately restore our system by replacing the main TriCaster unit. Procedural manuals and troubleshooting information are particularly important, and we are constantly adding to our knowledge base in this regard using a range of internal tools.

3.8 Final Thoughts

Can we really achieve our dream for the studio? It's one thing to look at cutting-edge technology and muse about whether something is or isn't possible from a technical standpoint, but it's also crucial to get those technologies working properly in real-world operations. Our professional studio operators have dedicated themselves to their craft for 20 or 30 years, so they work with real speed. Unwieldy and unresponsive systems thus tend to be shot down pretty quickly, so it will be important for us to figure out what level of production our system is capable of handling and what sort of projects will play to the system's characteristics as we work out what direction to take going forward.

IIJ Studio TOKYO—IIJ's IP studio—has only just started up, and we are learning the basics of video production while taking cutting-edge technologies on board, working to enhance the potential and stability of IP in the video production space, and testing methods of managing and monitoring systems, and the like as we continue to innovate. We already have talented operators joining IIJ and coming in on secondment from IIJ Engineering, and we are gathering information from external sources and working with consultants, so a community of people aligned with the goals of the studio is starting to form and things are going in a very positive direction.

Looking ahead, IIJ will continue to engage in projects that help create a world in which people and networks can connect across all sorts of industries, including the video industry.

Atsushi Sumita

Video Delivery Business Section, xSP System Services Department, Network Division, IIJ

Mr. Sumita previously worked as a sales engineer at a video trading company, which involved deploying and supporting non-linear editing systems for post-production and broadcasting stations.He joined IIJ in July 2019. He participated in ST 2110 remote production proof-of-concept and other projects, and he is involved in the basic design and operation of video production systems used in Olympics-related streaming and, currently, the system used to stream the Spring Festival in Tokyo.

The "3.2 Overview of IIJ Studio TOKYO Facilities" section was written by:

Ryota Imanishi

Video Delivery Business Section, xSP System Services Department, Network Division, IIJ

In 2015, Mr. Imanishi joined IIJ Engineering and became part of IIJ's video delivery business. His work includes the operation and maintenance of CDN services and the recording and streaming of events.

The IIJ Backbone—30 Years of Transformations

4.1 Introduction

The IIJ backbone, which started with just a few routers back in 1993, has now evolved into a large-scale network encompassing thousands of network devices. We have had to deal with a range of issues in the process due to technology and hardware being unable to keep pace with the rapid growth of the Internet—these issues have included communications and routing technologies, the limits of router hardware performance, and power supply issues. Looking back on these days, it was knowledge and ingenuity that got us through and enabled us to provide a stable Internet environment.

In the first half of this article, we discuss the background to and reasons for changes in the IIJ backbone over time as well as the innovations made, with some historical context mixed in. In the second half, we discuss the security measures IIJ has implemented in its network operations.

4.2 IIJ Backbone Through the Years

4.2.1 1993–2002: Early Years (Struggling with Resource Shortages)

Back when the Internet was edging toward a transition from academic to commercial use, people's awareness of the Internet was still low, system environments for connecting to the Internet left a lot to be desired[*1], and usage fees were high[*2], so the Internet was mainly being used on a trial basis.

Changes in Physical Configuration in the Early Years

The backbone started with a single configuration, with one backbone router installed for each POP (point of presence) and one dedicated line connecting the POPs in a daisy chain.

The IIJ backbone continued to grow with this single configuration, but once it became more and more common

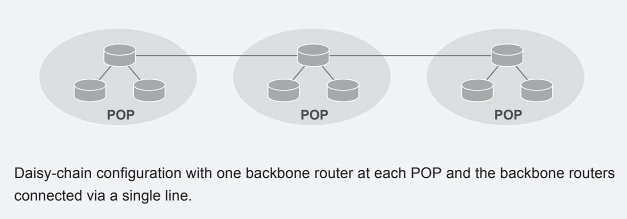
POP
POP
POP
Daisy-chain configuration with one backbone router at each POP and the backbone routers connected via a single line.

Figure 1: The Initial Backbone

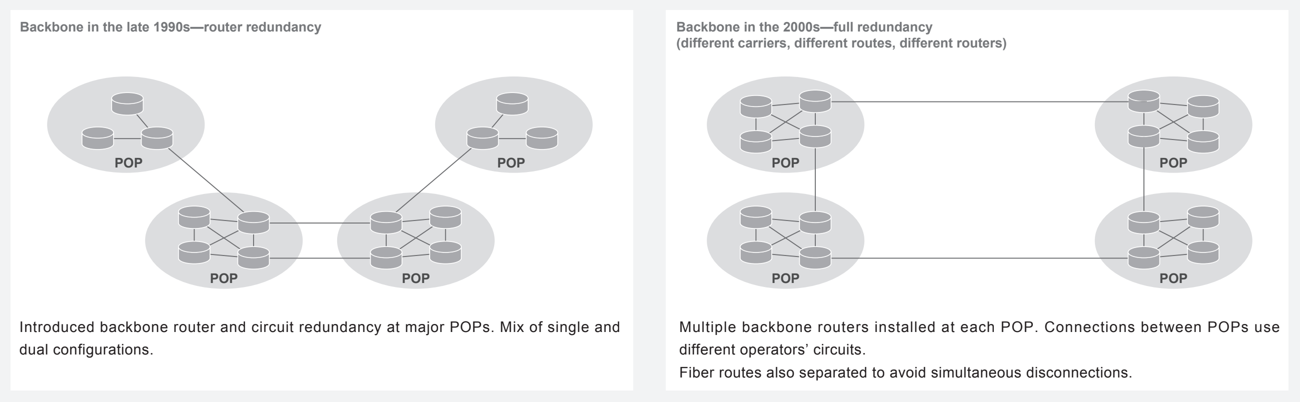
Backbone in the late 1990s—router redundancy
POP
POP
POP
POP
Introduced backbone router and circuit redundancy at major POPs. Mix of single and dual configurations.
Backbone in the 2000s—full redundancy
(different carriers, different routes, different routers)
POP
POP
POP
POP
Multiple backbone routers installed at each POP. Connections between POPs use different operators' circuits.
Fiber routes also separated to avoid simultaneous disconnections.

FIgure 2: Expansion of the Backbone

[*1] The TCP/IP stack became standard in Windows and Mac OSes around 1995.

[*2] Connecting to the Internet via a 45Mbps dedicated line cost 2,000 yen/month.

for companies to use the Internet for financial transactions and the like around 1999, quality demands on the Internet started to become more stringent, and this is when efforts to improve fault tolerance really started to ramp up. We started working on circuit and equipment redundancy at major POPs, and achieved full redundancy around 2002.

An instructive aside

In the winter of 2002, the Fukuoka POP, despite having two backbone circuits, ended up being isolated after both of those circuits were disconnected. This happened because water entered the fiber cables in a section of the fiber route that was common to both circuits, and then froze, causing damage. This lesson taught us to use separate routes for backbone circuits.

Changes in Routing Control During the Early Years

Since the beginning, we have continued to use BGP[*3] for EGP (Exterior Gateway Protocol) routing and OSPF (Open Shortest Path First) for IGP (Interior Gateway Protocol) routing.

In terms of the routing protocols used to control routes, we use EGP for user network routing information, such as customer and pool addresses, and we use IGP for routing information related to network configuration (devices and links between devices). This is also unchanged since the beginning.

IIJ backbone routing is designed with the idea of achieving a simple but robust network in mind, with the basic policy being to use EGP to propagate information on where networks are and IGP to control communication routes to the target networks. Initially, the number of routers and the total number of routes (full routes) on the Internet were small, BGP was still in development and thus only implemented the bare minimum of functionality, and we basically used a full-mesh iBGP configuration.

As awareness of the Internet began to grow globally, the IIJ backbone continued to expand, with the number of routers and routes increasing. As the number of neighbors and the amount of routing information sent and received increased, router restarts due to maintenance or failures put high demands on memory, and stability issues started to appear—for instance, it would take dozens of minutes or repeated restarts for routes to converge. The overseas routers, in particular, which are responsible for transmitting full routes to all routers in Japan, were coming up against their limits amid latency and the like. In the US, we had a pretty tough time as we went about procuring and increasing memory resources, and eventually we also brought in BGP

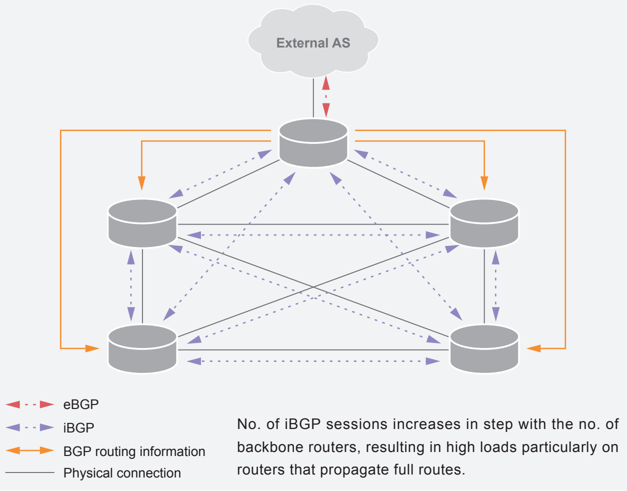

No. of iBGP sessions increases in step with the no. of backbone routers, resulting in high loads particularly on routers that propagate full routes.

Figure 3: iBGP Full Mesh Example

[*3] BGP (Border Gateway Protocol): Initially, BGP3 was the mainstream protocol. The implementation of CIDR and the like later led to BGP4 (RFC 1771). We have used BGP4 since beta testing.

route reflection (RFC1966[*4]) with the aim of reducing loads associated with sending and receiving routing information.

The first thing we did was create a three-cluster configuration spanning East Japan, West Japan, and overseas (Figures 4, 5, and 6).

The use of broadband connections spread and traffic volumes increased substantially from around 2001, and we continued to strengthen and expand the IIJ backbone. It was clear to us that we were fast approaching system limits, so we subdivided the clusters and shifted to a configuration in which we have a cluster at each POP (Figure 7). And more than 15 years on, the backbone is still based on this configuration.

A memorable aside

At that time, the sheer growth in the number of routes was a serious problem. In our chassis-based routers, even our modular interface cards were on the verge of running out of memory, and we had to work late into the night to add memory to hundreds of cards to avert system malfunctions.

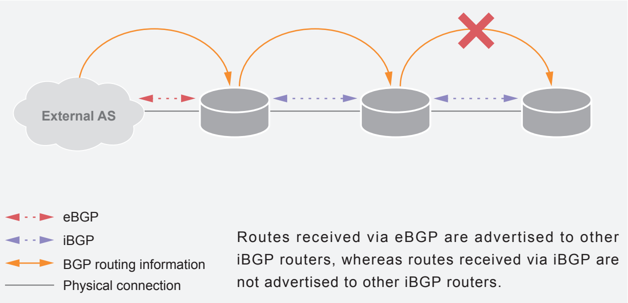
External AS
eBGP
iBGP
BGP routing information
Physical connection
Routes received via eBGP are advertised to other iBGP routers, whereas routes received via iBGP are not advertised to other iBGP routers.

Figure 4: Normal BGP Adjacency Diagram

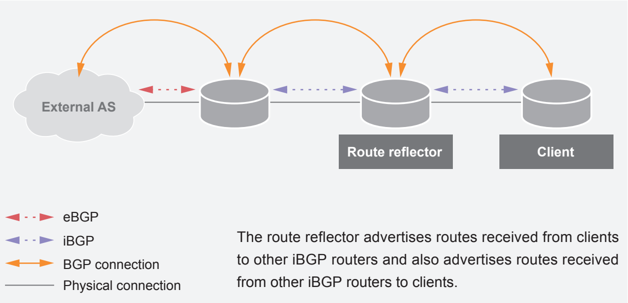
External AS
Route reflector
Client
eBGP
iBGP
BGP connection
Physical connection
The route reflector advertises routes received from clients to other iBGP routers and also advertises routes received from other iBGP routers to clients.

Figure 5: RR-RC Adjacency Diagram

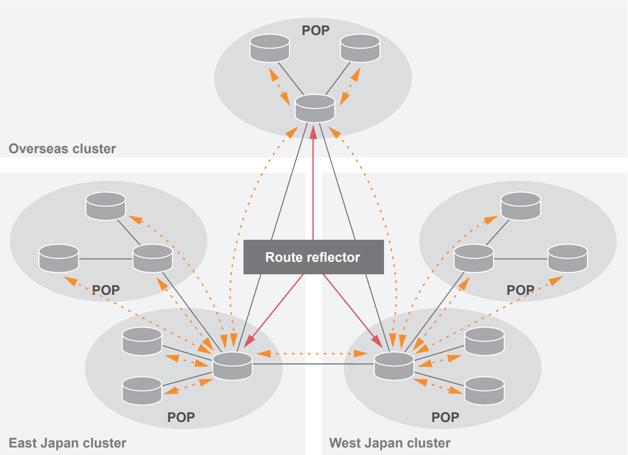

Figure 6: Overview of the Clusters

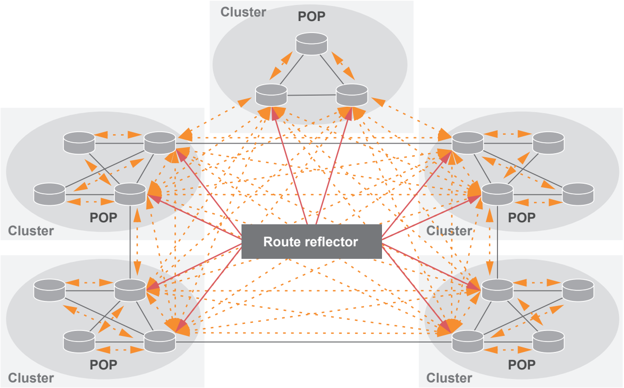

Figure 7: Overview of Cluster Subdivisions

[*4] RFC 1966 BGP Route Reflection—An alternative to full mesh IBGP.

A simple aside

It used to be whispered in the industry that only around 50 routers could coexist in any one OSPF backbone area because of the CPUs being underpowered. In light of this, we split up the OSPF areas on the IIJ backbone too, which dramatically increased the operational difficulty level, and we even experienced a number of accidents because of this. Fortunately, hardware subsequently evolved, and we were able to do away with the area divisions, but this is a prime example of why a simple configuration is best.

The Internet is an amalgamation of networks managed and operated by many different entities, and unintended route hijacking due to misoperations does occur on rare occasions. This can, in some cases, impact the entire Internet, so considerable care must be taken with respect to the routing information sent to and received from other ASes. At IIJ, we introduced route filtering on all edge routers (including within IIJ itself) early on to prevent users from becoming the cause of problems.

Initially, we did this using a combination of access control lists and AS path filtering, but the route filters on all edge routers need to be updated every time you added a new CIDR block, provider-independent address, or the like, so it was very complicated and a lot of work. So we decided to adopt BGP Communities Attribute (RFC 1997). It is quite simple to use. You add a BGP community at route inflow sources and route origins, propagate this inside the backbone, and control route advertising on the edge routers based on the BGP community. This made routing control much easier and greatly reduced the number of settings to be changed, helping to stabilize operations.

To recap, the various technologies were still being developed in these early years, and struggled to keep up with the Internet's growth. This was an era of working to solidify our foundations through trial and error while constantly battling resource shortcomings.

4.2.2 2003–2006: Popularization Era (Rising Quality and the IPv6 Rollout)

As use of the Internet spread, so did demands for quality. We progressively introduced redundancy into the backbone from around 2000, and although lengthy interruptions mostly disappeared, packet losses due to route changes started to become an issue.

Dynamic routing protocols like BGP and OSPF are used for Internet routing control, enabling automatic rerouting of communications in the event of failures. With this sort of dynamic routing, changes in the network are propagated throughout the network as routing information, and each router receiving the information creates its own routing table, ensuring that the entire network can function normally without any inconsistencies. The convergence of state changes resulting from this series of operations is called routing convergence, and the time taken to reach convergence (convergence time) is one measure of network quality and performance.

The state changes leading up to convergence can be roughly divided into the following phases.

- **Event detection (router addition/deletion, link up/down, configuration change, etc.)**
- **Injection into routing protocols**
- **Propagation of routing information**
- **Routing calculations (for each routing protocol)**
- **Incorporation into the routing table**

State changes occur frequently on the Internet due to maintenance and failures. When a state change occurs, convergence needs to be reestablished, and while this is happening, inconsistencies between different routers' routing tables can cause packet losses. The larger the network, the longer convergence times tend to be, and the greater the impact of convergence performance on network quality. So speeding up routing convergence is crucial to achieving a more stable, higher-quality network.

Technologies for speeding up routing convergence were just beginning to emerge at the time (circa 2003). We first decided to study IIJ's backbone performance, measuring it using equipment scheduled to be decommissioned. We checked the results against device debug logs to determine what was taking up time, and we then looked at potential countermeasures. We ended up taking the following three major actions.

- **Router upgrades**
- **Parameter tuning**
- **Switch to topology that makes it easier to detect outages**

It was impossible to tune the various parameters unless the routers were upgraded to the latest OS. We needed just under a year for the backbone routers alone, and several years to complete this on all routers. Alongside this upgrade, we also added IPv6 support. Starting with the upgraded routers, we set about tuning parameters and shifting to a dual-stack network. BFD (Bidirectional Forwarding Detection) was not yet available at the time, so we did the best we could, changing to point-to-point on L2 segments to the extent possible and implementing a topology that would not rely on keep-alive. The upshot of our efforts was that we achieved a convergence time of under a second.

Alongside our efforts to speed up routing convergence, we also worked on the development of a range of systems to improve quality.

- **System for monitoring state changes based on router logs**
- **System for recording routing updates**
- **System for measuring and monitoring packet losses and delays between points**

The sort of quality we achieved is taken for granted today, but it was through these efforts that we achieved it at an early point in time..
To recap, the various technologies were still being developed in these early years, and struggled to keep up with the Internet's growth. This was an era of working to solidify our foundations through trial and error while constantly battling resource shortcomings.

4.2.3 2007–2010: Battling with Traffic (Shift to BF Routers)

With traffic ever increasing, our routers were coming up against their limits, so in line with the design at the time, we considered using OC-768 (40G) as our next connection media after OC-192 (9.6G). Of the routers that met our requirements, only the Cisco CRS-1 supported OC-768. But we quickly gave up on that idea as it had a one-ton floor load and we could not install it. Hence, our only option was to add multiple 10GbE, and so we faced the need to rethink our design given issues with the routers' 10GbE port count and capacity.

Our solution was to use multiple routers to implement the 3-stage switch fabric architecture used to increase the capacity of the CRS-1 backplane, thus creating a giant virtual router out of a "router group" (Figure 8-1).

With this concept, the backbone routers (denoted BF) corresponding to the switch fabric must be connected to all the backbone routers (denoted BB) at the edge. Looking at it from the other end, however, the BB routers need to have as many ports as there are BF routers, but because they handle incoming/outgoing at the edge, you can't use all of the ports to connect to the BF routers. Based on what capacity we would need if the capacity of the router group were to double every year for four years (16 times current traffic), we calculated how many ports could be used to connect to the BF routers out of the BB routers' maximum port count.

Having solved the connection port issue, we then took advantage of the fact that we were using multiple routers and worked on the idea of distributing them across multiple locations, instead of having them all in the same place, to reduce the overall number of routers. We looked at distributing the BF routers across three locations in Tokyo where traffic was heavy (Figure 8-2). The problem with distributing the routers like this is that you need a huge number of 10G lines between sites. Given the unit cost of 10GbE circuits at the time, we surmised there would be a hefty price tag, such that it would be cheaper not to use a distributed deployment. This led us to speak to communications carriers about the number of circuits we envisioned and what the price per circuit would be. We compared this with what the per-circuit cost would be if we were to operate transmission equipment ourselves at scale, and we decided to have our own transmission equipment on some sections. While we had been using simple transmission equipment, full-fledged equipment presented both a high hurdle and a high price tag. But we spent some time approaching manufacturers, testing their equipment, asking questions, and having them explain things. Sensing how earnest we were, one of those manufacturers decided to work with us at the price level we were hoping for, and this ended up being a deciding factor.

With the circuits for our distributed deployment sorted out, we set about designing the circuits between the router group and each of our locations. A conventional 1+1 redundancy design would require a huge number of circuits, so we looked at N+1 redundancy to reduce the cost, but figuring out how to distribute things with an odd number of circuits was extremely difficult, and we couldn't find a good method for this (Figure 8-3). In the end, we gave up on part of the 3-stage fabric concept and decided to implement N+1 redundancy by connecting the BF routers between major locations like Tokyo and Osaka. Because connections run through the BF routers, without a clear picture of which Tokyo/Osaka circuit traffic coming in from the BB routers is using, and what the volume of that traffic is, we would not be able to properly plan for capacity expansions or traffic rerouting during outages and maintenance, and this creates a very difficult problem. NetFlow analysis is the only way to solve this problem, and if that analysis is time consuming, you can't cope with sudden traffic spikes. It happened that right around that time we were developing a system for high-speed analysis of distributed systems*5, and this helped us avoid creating congestion during outages and maintenance. We designed our system to last four years, but we ended up doubling that as we were able to continue expanding it, without any changes to the design, for eight years.

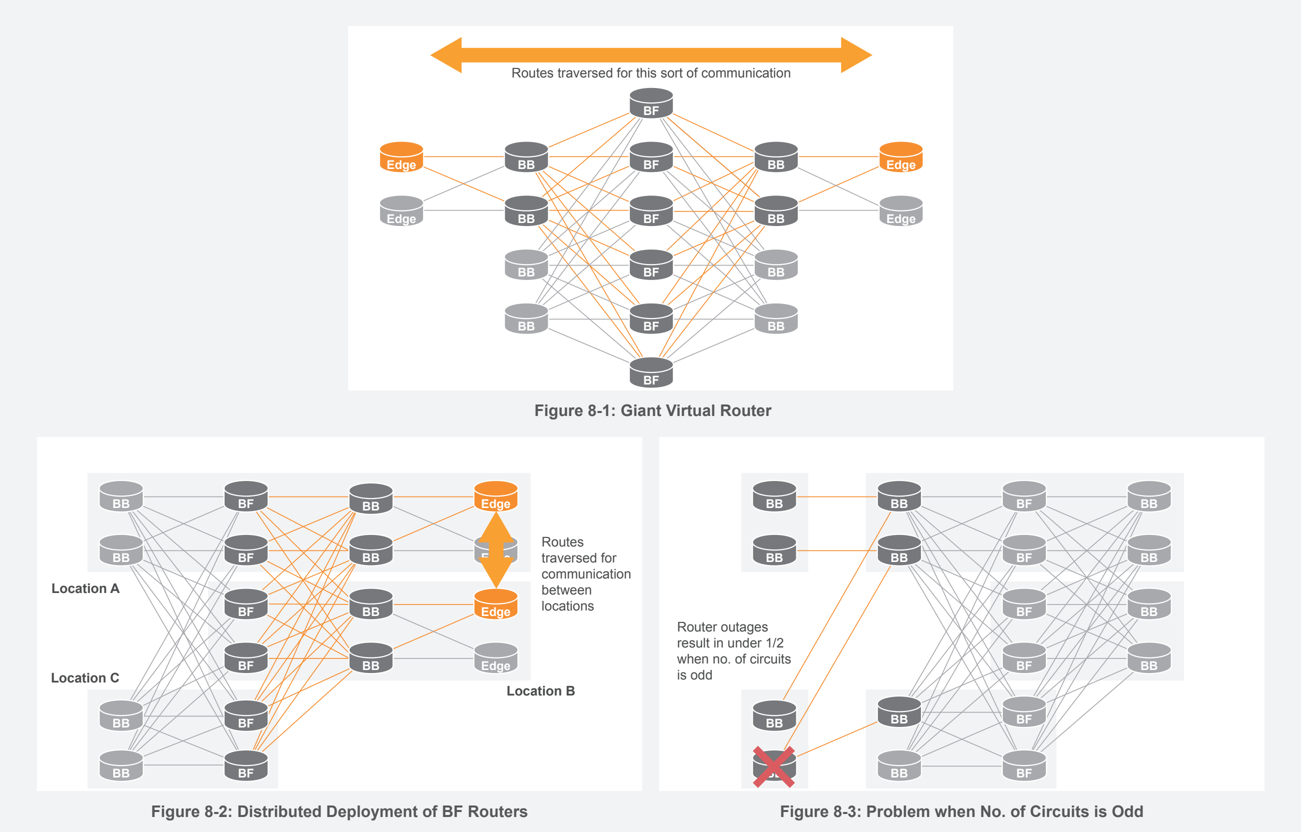
Routes traversed for this sort of communication
Edge
BB
BF
Edge
BB
BF
Location A
Location C
Location B
Routes traversed for communication between locations
Router outages result in under 1/2 when no. of circuits is odd

Figure 8-1: Giant Virtual Router

Figure 8-2: Distributed Deployment of BF Routers

Figure 8-3: Problem when No. of Circuits is Odd

Figure 8: Overview of Internet Backbone with Fabric Configuration

*5 Refer to Chapter 3, Cloud Computing Technology "Implementation and Application of the DDD Distributed System" in IIR Vol. 4 (https://www.iij.ad.jp/en/dev/iir/004.html).

4.2.4 2011 Onward: Network Cloud (Building an Integrated Core and Expanding Private Areas)

Cloud services were going into full swing around this time, with AWS, GCP, Azure, and the rest already on the scene, and IIJ had also released IIJ GIO. Demand for communications between private sites isolated from Internet traffic was growing in conjunction with this, and we used MPLS/L3VPN to expand our private backbone separate from the Internet backbone.

The Internet backbone uses a fabric configuration as described above, and it was designed to be capable of transporting overall traffic as the fabric routers were scaled out, but as we only had 10G media, the more traffic grew, the more operational issues we encountered. For connections within POPs, we used load-balancing methods such as LAG (Link Aggregation) and IGP/BGP multipath, but there are limits to how well you can distribute traffic with traffic flow hashes, and you end up consuming too many ports. You also don't know which links IP packets are flowing through, so it's difficult to confirm that the system is running normally, and thus we were very much looking forward to consolidating everything on 100G.

We actually started using 100G around 2012, opening a 100G connection to JPNAP. Our Internet backbone was already designed with fault tolerance in mind between Tokyo, Nagoya, and Osaka—we used around 20 OC-192 circuits spanning different carriers, different routes, and different locations. Simply switching to 100G with this configuration would be too costly, so we took advantage of the capacity offered by 100G to achieve the following without losing any redundancy in terms of routes or locational distribution.

- **Integrated different locations' backbone network circuits**
- **Eliminated Tokyo/Osaka-dependent structure**

Since IIJ does not have its own fiber, we had to procure carrier circuits for long-distance sections. So that we would not have to obtain separate carrier circuits for each of the multiple network planes, we enabled MPLS/L2VPN pseudowires (PW), allowing bandwidth to be shared by multiple network planes. Further, because maintaining route and locational distribution redundancy for each network would increase the operational effort and costs involved, L2VPN handles most of the route distribution and traffic engineering, and MPLS high-speed rerouting conceals topology changes due to circuit failures and the like, and this configuration makes it easier to control each network. On our high-traffic Internet backbone, we shifted to a simpler configuration whereby core POPs are fully meshed, eliminating transit traffic between core locations.

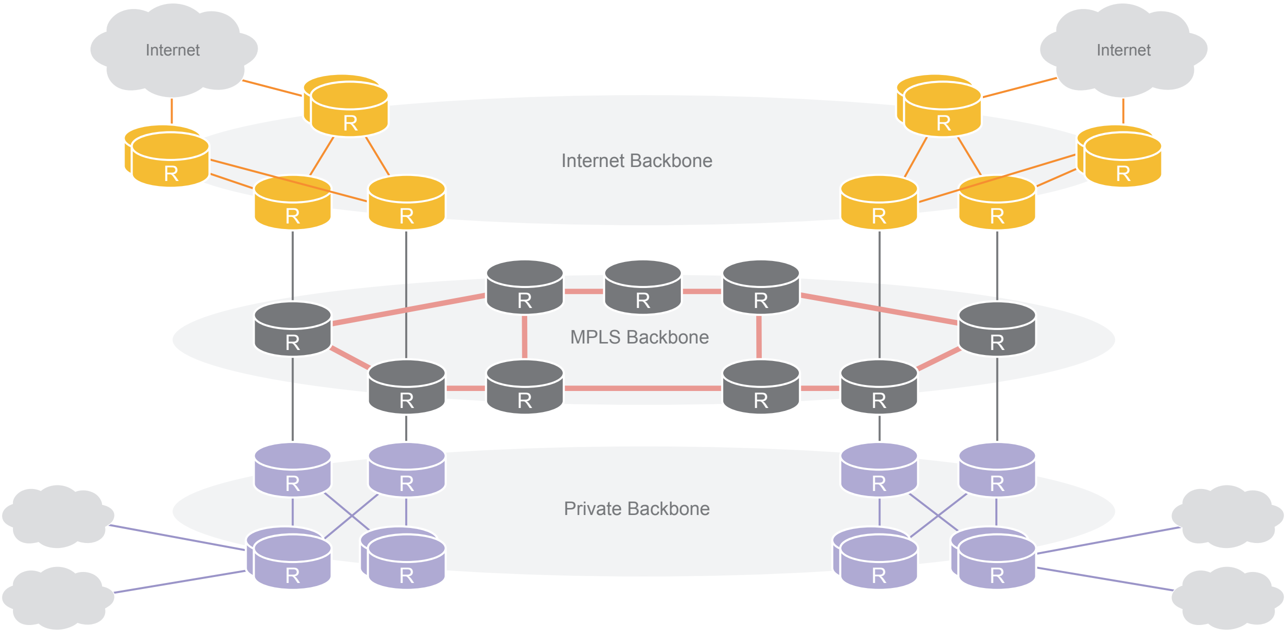

Figure 9: Overview of 100G Core Backbone

Before we eliminated the Tokyo/Osaka-dependent structure, the backbone was configured so that even locations outside the Tokyo/Osaka vicinity were tied to Tokyo or Osaka respectively. That was fine from the perspective of traffic efficiency, but it also meant that communications would go out in those non-central locations in the event of disasters affecting Tokyo or Osaka. We have spent several years addressing this. To increase fault tolerance, we extended our Sapporo and Sendai circuits to Nagoya via non-Kanto (i.e., non-Tokyo) routes, and our Okayama, Hiroshima, Fukuoka, Matsue, and other circuits to Nagoya via non-Osaka routes. We have been working on increasing fault tolerance for several years, and our international circuits between Japan and the US are distributed across Tokyo, Osaka, and Nagoya.

IIJ's network originally only had a location in the US, the central hub of Internet traffic, but alongside these efforts, we also extended our network to other regions—Europe in 2013, and Hong Kong and Singapore in 2014. This means that traffic can now be exchanged directly with Asia and Europe, so we are not reliant solely on the US for international connectivity, and this has improved Internet connectivity.

Implementing an integrated 100G backbone like this meant we could also smoothly expand our private backbone, which is small in comparison with Internet traffic. We have progressively expanded as a cloud exchange to facilitate interconnectivity with public clouds, and we have expanded as a network cloud to meet the needs of today's increasingly diverse workstyles.

We began this section with the IIJ backbone's early years, following its history up to the present and looking at the repeated improvements and changes made over the years to address the prevailing issues of the day. The Internet is an integral part of society's infrastructure, and people will no doubt demand even greater levels of reliability going forward. IIJ will continue to expand its systems to provide reliable social infrastructure that serves the needs of people everywhere.

4.3 IIJ's Network Security Measures

IIJ has also worked to improve security to ensure its networks can be used appropriately. In this section we look at some of the measures IIJ has taken with respect to the security of network operations.

4.3.1 Source Address Validation

On the Internet, routing information for delivering IP packets to their destination is basically searched for based on the destination IP address given in the IP packet header. The IP packet header also contains information on the source IP address, but the IP packets will be delivered to the destination even if this information is incorrect. Upon receiving the IP packet, the destination determines where it came from by looking at the source IP address in the IP packet header and, if necessary, sends out a response packet. If the source IP address is wrong, the system will still take the incorrect IP address information to be true and send the response packet to entirely the wrong host. This behavior is exploited by malicious attackers, and a variety of attack methods have been devised and used in real-world attacks. Attackers can use these methods, for example, to make it difficult to identify the source of an attack, to hijack communications by spoofing another host, or to have response packets sent to a specific host.

DNS reflection attacks (DNS amplification attacks) exploiting the DNS system have been observed since around 2005. These attacks involve spoofing the source IP address to be the IP address of the attack target and sending DNS queries to nameservers as a stepping stone. The nameservers respond, with the name resolution resulting in an increased amount of data, efficiently exhausting the attack target's bandwidth, the aim being to disable service. Attackers hijack Internet-connected hosts in advance and then carry out such attacks by sending packets with spoofed source IP addresses from those hosts. To ensure that IIJ's connectivity services are not exploited in such attacks, we decided to introduce technology that prevents source IP address spoofing.

Problems associated with IP spoofing were recognized early on, with certain problems and countermeasures being documented in RFC 2827 (BCP38) and RFC 3704 (BCP84). To combat this, you need to verify whether an appropriate source IP address is used as close as possible to where the connection service is terminated. The methods of source address validation available on the equipment IIJ was using at the time were unicast reverse path forwarding (uRPF), which uses a route search mechanism, and access control lists (ACLs), which use packet filtering. We implemented these as appropriate given the functional limitations of the different equipment models and software versions. In March 2006, we announced[*6] a rollout of sender verification on all connection services, which we subsequently completed. This has prevented IIJ's connection services from being exploited in attacks, improving the security and facilitating the stable operation of the overall network.

4.3.2 Internet Routing Registry (IRR)

With a variety of different networks connecting to the Internet and expanding, one major consideration is how to go about coordinating BGP routing control policies among networks. To address this, the IRR publishes routing policies as objects and provides functionality allowing network operators to query each other's policies. Objects registered in the IRR database can be used to automatically generate route filters, perform checks when failures occur, etc. At IIJ, we register the main types of objects commonly looked up in the IRR database—such as route, route6, and as-set—and keep the information up to date. We have been using Merit RADb, a service run by Merit since the 1990s, to register our IRR routing objects. Alongside this, since 2005 we have also used the JPIRR operated by JPNIC, and at present, we mainly use these two IRRs.

If objects registered in the IRR database are rewritten without permission, other networks that look those objects up could generate the incorrect route filters, which could cause reachability problems for IIJ. With both Merit RADb and JPIRR, objects are basically updated by sending emails to the administration system. There are several authentication options available when doing this. At IIJ, we use the strongest one: Pretty Good Privacy (PGP). This sort of authentication involves verifying PGP digital signatures. To use it, you first register a PGP public key with object modification permission in the IRR database. IIJ completed the transition to PGP authentication in 2003.

4.3.3 Resource Public Key Infrastructure (RPKI)

RPKI is a public key infrastructure for certifying the distribution of number resources such as IP addresses and AS numbers, and using RPKI can help improve routing control security. An organization that receives an IP address can issue a Route Origination Authorization (ROA) from the RPKI system to indicate which AS should be advertising the network. Using this information, it is possible to verify whether route information received via BGP was generated by a valid origin AS. IRR is more widely used at present, but RPKI is better for automation and can use more reliable information, so we expect the use of RPKI to spread.

IIJ completed the issuance of basic ROAs in 2020. The ROAs include a maximum prefix length indicating the extent to which the AS can split routes it advertises, but since IIJ does not split advertisements, we leave the prefix length of advertised routes as is. This is also what is recommended in RFC 9319 (BCP185). Also in 2020, we introduced a policy for using ROA information to verify BGP route advertisements received from peers and upstream and discarding route information that is inconsistent with the ROAs. On the IIJ network, this makes it possible to identify and discard routes even if the ROA-issuing network receives a route advertisement from an incorrect origin AS. Merit RADb also implements a feature that automatically removes objects that are inconsistent with ROAs,

[*6] IIJ to Roll Out Source Address Validation on All Connection Services (https://www.iij.ad.jp/news/pressrelease/2006/pdf/0308.pdf, in Japanese).

so issuing ROAs is also a way to avoid registering incorrect objects in the IRR database.

4.3.4 Mutually Agreed Norms for Routing Security (MANRS)

Network security measures used in the operation of the Internet become more useful when a large number of network operators adopt them in concert with one another. MANRS is a voluntary global initiative promoting the introduction of such measures. With support from the Internet Society (ISOC), MANRS sets out recommended security measures (actions) for different areas, which it asks organizations involved in the Internet's operation to put into practice. Organizations that approve of this approach can become a participant by informing MANRS of the actions they have implemented.

IIJ appropriately implements security measures suited to its operations. These are consistent with the practices recommended by MANRS, and IIJ joined MANRS in 2015 as the first participant from Japan[*7]. Looking ahead, IIJ will continue to review its operations and continuously make improvements to ensure the stable operation of the Internet.

1993–2002: Early Years (Struggling with Resource Shortages)
Toshio Iwasaki
Manager, Operation Technology Department, Infrastructure Engineering Division, IIJ

2003–2006: Popularization Era (Rising Quality and the IPv6 Rollout)
Yoshio Asano
Infrastructure Technology Department, Network Division, IIJ

2007–2010: Battling with Traffic (Shift to BF Routers)
Kunio Kataoka
Infrastructure Engineering Division, IIJ

2011 Onward: Network Cloud (Building an Integrated Core and Expanding Private Areas)
Fumiaki Tsutsuji
Network Planning Manager, Network Technology Department, Infrastructure Engineering Division, IIJ

IIJ's Network Security Measures
Yoshinobu Matsuzaki
Technology Development Section, Operation Technology Department, Infrastructure Engineering Division, IIJ

[*7] MANRS Turns 1 and First Japanese Operator, IIJ, Joins (https://www.manrs.org/2015/11/manrs-turns-1-and-first-japanese-operator-iij-joins/).

May 2023 Vol.58

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.

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/