IIR

Internet
Infrastructure
Review

July 2021
Vol. 51

Periodic Observation Report

FY2020's 200x Rise in Spam, and Tackling Password-protected ZIP Files

Focused Research (1)

IIJ's Further Challenges as a Full MVNO —Road to commercial NSA/SA services through a unique path in local 5G environments

Focused Research (2)

Meet Barry, IIJ's Tool for Rapid Fault Resolution

IIJ
Internet Initiative Japan

Internet Infrastructure Review

July 2021 Vol.51

Executive Summary

I'm writing this executive summary in Tokyo less than two months before the Tokyo Olympics are set to commence after being postponed for a year. The Tokyo Olympics will feature 339 events representing 33 different sports, with over 10,000 athletes likely to compete. The number of spectators and event staff was expected to exceed 10 million, and while this figure may be heavily reduced under current circumstances, the Olympics remain a rare global spectacle and will still be a truly special competition for the athletes.

Information and communications technology (ICT) is a key supporting element behind huge events like this. The smooth running of the Games rests on a whole range of systems that provide support including measuring athletes' performances, conveying event results to media outlets and the like, authentication and venue admittance for athletes, staff, and spectators, and the online distribution of data and imagery that provide a vivid picture of the athletes' endeavors. The athletes give amazing performances that are broadcast around the world, bringing excitement to people everywhere, and all the while ICT is making a huge contribution behind the scenes.

Although the impact of COVID-19 prevents the Games from being held in their originally conceived form, everyone involved is still no doubt hard at work preparing for the event. And so too in the case of ICT. With people's movements restricted across the globe, ICT continues to help people connect with one another and enjoy their leisure time. My hope is that the Olympics can be held safely under these difficult circumstances, and that as many people as possible are able to enjoy the athletic performances and excitement via 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.

In our periodic observation report in Chapter 1, we summarize trends in the messaging space over the past year, with a focus on email. One notable finding comes from the analysis of spam arriving at honeypots run by IIJ during the first half of the previous fiscal year (April 2020 – March 2021), which shows an unprecedentedly large volume of spam being received. Other topics covered include the peculiar Japanese practice of attaching encrypted ZIP files to emails and the use of online conferencing systems for phishing, which created a stir last fiscal year.

The first focused research report in Chapter 2 discusses IIJ's efforts in the area of 5G NSA (non-standalone) and SA (standalone). In the 4G space, IIJ provides services as a full MVNO with some of the features of a core network. 5G NSA involves making enhancements to the 4G core network to provide high-speed 5G communications. With 5G SA, a new 5G core network is used to provide 5G services. In the case of both NSA and SA, we have tested the technologies in-house and performed proof-of-concept work at our Shiroi Wireless Campus testbed, and we are using some of the results gleaned from this work to provide regional BWA and local 5G in the cable television industry.

The second focused research report in Chapter 3 introduces Barry, an in-house system we use when dealing with system faults at IIJ. As systems become larger and more complex, demands for reliability are only increasing, and whenever a system incident occurs, it is extremely important for IIJ, as a provider of ICT services, to detect the issue, provide accurate information to the necessary personnel, and swiftly enact a response. Instead of relying on third-party tools to support this incident response process, we made the decision to develop a system in-house with an eye to improving our own workflow and creating new technologies. The report also describes the background to Barry's development, which was informed by feedback from engineers actually involved in IIJ's operations.

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

Junichi Shimagami

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

FY2020's 200x Rise in Spam, and Tackling Password-protected ZIP Files

1.1 Introduction

A year has passed since the sudden rise of telework, from both home and satellite offices, as a means of dealing with COVID-19. While the world has been transformed over the past year, the importance of email for enterprises remains unchanged.

In this article, we look back on events in the messaging space in 2020.

1.2 FY2020 Spam and Virus Data

Figure 1 shows the number of spam emails received by honeypots run by IIJ over April 2020 – March 2021.

In the first half of fiscal 2020, we intermittently observed unprecedented levels of spam. It is difficult to see because the vertical axis accommodates the maximum value, but the average for April is set to 1. In the first wave in early May, we received 10 times [1] as many spam emails, the second wave in mid-May brought around 40 times [2], and around 60 times as many came in late May through early June. This continued intermittently, with around 200 times as many spam emails received at end-July. I think it's rare for the equipment designs of organizations in general to call for investment in equipment capable of withstanding 200 times normal levels, so calling this a DDoS attack would not really be an exaggeration.

Let's also look at the virus figures. Figure 2 shows the total number of viruses arriving at IIJ's honeypots during the same period.

The average for April is again set to 1 here, and 1,000 times that many viruses were received in June. But the other figures are hard to see because the range of counts observed over the year is too wide, so Figure 3 plots the same data logarithmically.

There is always something going on, but unlike with spam, incoming virus volumes are concentrated over short time intervals several times a year.

A look at a sample email received in June reveals the subject line "Look at this photo!", and the name of the attached ZIP file—IMG135123.jpg.js.zip—makes you think it could contain an image. Upon unzipping, we find it contains JavaScript that downloads malware. The file itself is not malicious per se, but if the user runs the JavaScript, it will download the malware. Getting a user to open a file by making them think it's an image is a classic old technique for those in the know, but we must take note of this and not regard it as an antiquated attack.

The second largest peak is in September. The virus observed on this occasion was Emotet. A characteristic of

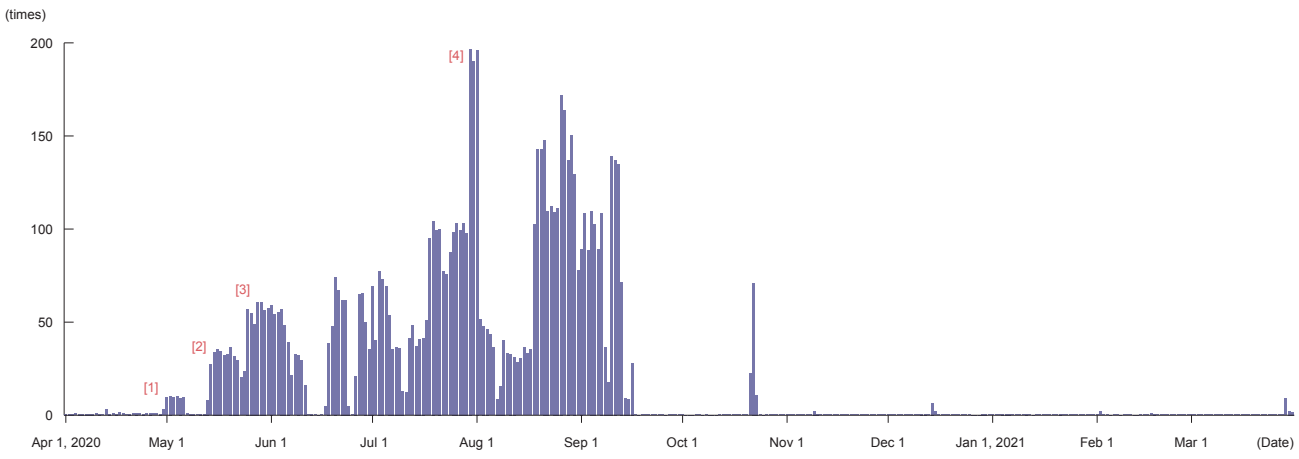

Figure 1: Spam Arriving at IIJ Honeypots (Apr 2020 – Mar 2021)

Emotet, which ran rampant in 2020, is that it encrypts itself in a password-protected ZIP file, but the observations in September revealed a previously unseen characteristic. The output in Figure 4 is part of a sample showing the detailed structure of the Emotet received (encrypted in ZIP format).

The encryption method had changed from ZipCrypto, the ZIP file standard, to AES 256-bit.

AES offers stronger encryption than the ZIP standard and is supported by some archivers such as 7-Zip, but Windows

```
$ zipdetails 'OXEFVNG1 20209月16.zip'

0000 LOCAL HEADER #1       04034B50
0004 Extract Zip Spec      33 '5.1'
0005 Extract OS            00 'MS-DOS'
0006 General Purpose Flag  0803
     [Bit  0]              1 'Encryption'
     [Bit 11]              1 'Language Encoding'
0008 Compression Method    0063 'AES Encryption'
003D   Encryption Strength 03 '256-bit encryption key'
```

Figure 4: Part of Sample Showing Detailed Structure of Received Emotet (encrypted in ZIP format)

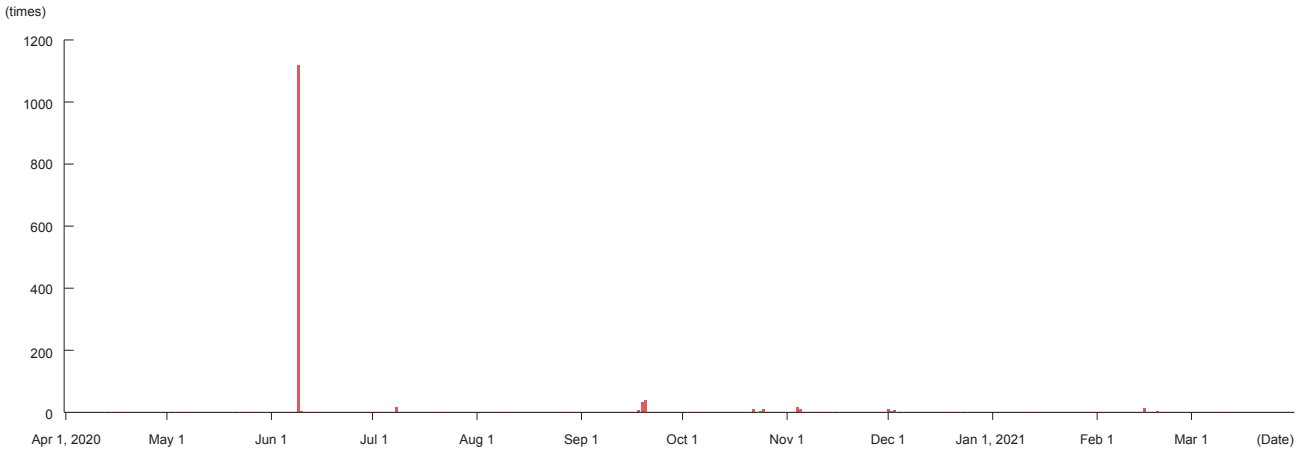

Figure 2: Viruses Arriving at IIJ Honeypots (Apr 2020 – Mar 2021)

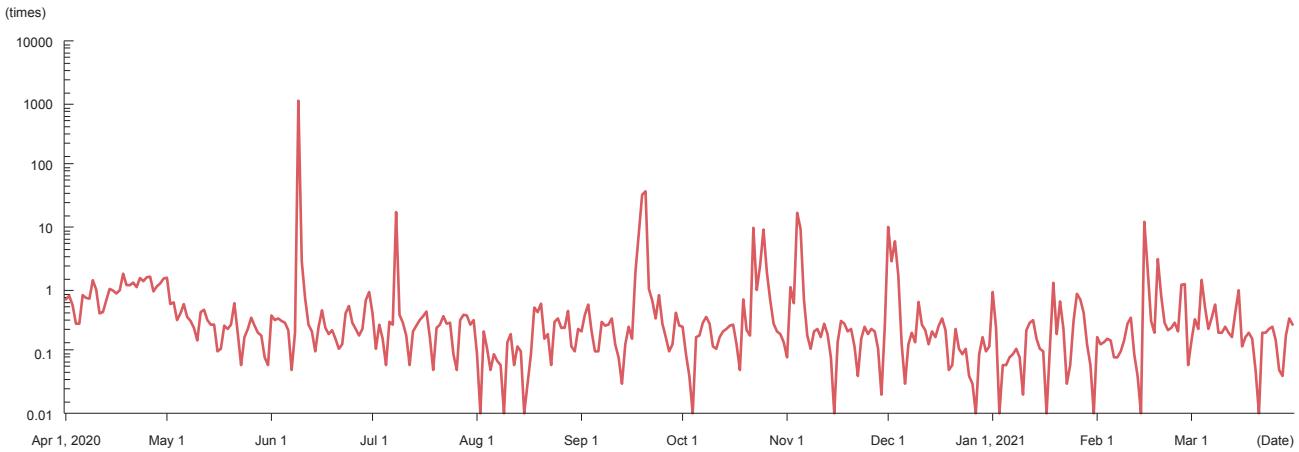

Figure 3: Viruses Arriving at IIJ Honeypots, Plotted Logarithmically (Apr 2020 – Mar 2021)

Compressed Folders do not support it. There would not seem to be much point in strengthening the encryption if your objective is to spread a virus, because you want as many targets to open it as possible. So why was AES being used?

This is just a guess, but it may be an attempt to avoid sandboxed detection. Some security products and services pull text strings likely to be passwords out of the email body and use them to try to open protected files in a sandbox. But the Windows standard does not support AES-encrypted ZIP files. Even if the correct password can be extracted from the email, the archive cannot be opened, so malicious behavior goes undetected and the file circumvents the sandbox protection.

So evidently, attacks that may appear the same on the surface are gradually changing tactics.

1.3 Calls to Abandon Encrypted ZIP Files

In November 2020, Takuya Hirai, Japan's Minister for Digital Transformation, spoke about the idea of "abandoning the use of encrypted ZIP files (a practice commonly referred to as PPAP in Japan)"[*1].

The phrase "encrypted ZIP files" here is referring to the practice of encrypting files in a password-protected ZIP archive and sending them to someone via email, and then sending the password in a separate email[*2]. The practice is somewhat unique to Japan and virtually nonexistent overseas.

Reasons cited for using this method include to prevent files being missent and for route encryption, and it is often stipulated mainly in organizations' IT security policies because it can be performed by individual users and the cost is low (the immediate cost at least, since no special applications are needed).

1.3.1 Problems with Encrypted ZIP Files

Why should encrypted ZIP files be abandoned? The answer is that the many risks outweigh the potential benefits of this method.

The most serious issue is that it prevents received emails from being scanned for viruses. As explained in Section 1.2, emails are still a useful attack vector for would-be attackers since anyone can send and receive them[*3]. So generally, many organizations implement virus protection on their gateways, Web servers, and the like to prevent infection by viruses attached to emails.

But you can easily circumvent such protection by encrypting attachments. This renders the antivirus product incapable of

*1 Cabinet Office, "Summary of press conference by Minister of State Takuya Hirai, November 17, 2020" (https://www.cao.go.jp/minister/2009_t_hirai/kaiken/20201117kaiken.html, in Japanese).

*2 At the time of publication, some emails sent by IIJ staff use this method. While work and customer circumstances mean we cannot completely abolish the practice immediately, we do intend to progressively discontinue encrypted ZIP files as preparations fall into place.

*3 Information security site for the Japanese public (Ministry of Internal Affairs and Communications), "Virus infection routes" (https://www.soumu.go.jp/main_sosiki/joho_tsusin/security/basic/risk/02-1.html, in Japanese).

analyzing the file contents[*4]. And this is perhaps an obvious outcome since the whole point of encrypting files is so that others can't read the contents.

Since around 2012, we have observed targeted attacks[*5] directed at specific organizations, individuals, and teams apparently designed to exploit this weakness. These attacks involved exchanging harmless emails on a number of occasions to gain the recipient's trust, culminating in a final email saying, "I've put all my questions into the attached file". The attached file would contain malware encrypted in ZIP format so as to deliver it to the recipient without a virus scan.

The Emotet version seen in 2020 would encrypt itself in a ZIP archive as a Word or Excel file with macros, which it attached to emails. To spread itself, the virus would read the address book and emails on infected devices and send out emails purporting to be replies coming from the infected device.

And Emotet is very sophisticated in that it comes in so many variants, spreading itself with email subject lines like "invoice" and "order" (in Japanese to target Japanese users) and timing its emails for the end/start of the month. Encrypted ZIP files are exchanged so commonly in Japan that unsuspecting users are likely to open them unless they are paying careful attention. This happening over and over again is probably what precipitated the virus's explosive growth.

So this method, intended by email senders to be a security measure, results in the recipient's email system unconditionally treating files as harmless, with the result that it actually exposes the recipient to risk.

The US CISA (Cybersecurity & Infrastructure Security Agency) has issued an alert on Emotet[*6]. The text states that blocking encrypted ZIP files that cannot be scanned by antivirus software is a risk mitigation measure.

1.3.2 Alternatives to Encrypted ZIP Files?

How should we send and receive files if we abolish the encrypted ZIP method?

There are moves within the government to use shared storage systems[*7]. Instead of sending an attachment, you send the path or a one-time URL to the file in shared storage.

This method has a number of advantages. Files can be scanned for viruses in the shared storage rather than in email[*8], missent files can be deleted as soon as the error is realized, and files that are too large to be sent via email can also be exchanged.

*4 Encrypted files are sometimes identified as viruses even when the contents cannot be analyzed. But because this is based on limited information, such as the file hash and filename, it is more difficult to identify viruses this way than with unencrypted files.

*5 Attacks that target a specific team or individual within an organization to allow an attacker to steal confidential information or gain a path to intrusion. While in technical terms such attacks do not differ from conventional methods such as sending people viruses or directing them to phishing sites, the emails do not end up in other organizations' or antivirus vendors' honeypots, so the attacks are harder to detect.

*6 CISA, "Alert (AA20-280A) Emotet Malware" (https://us-cert.cisa.gov/ncas/alerts/aa20-280a).

*7 When the Cabinet Office announced the policy, it did not disclose what sort of system would be used. But, unfortunately, a subsequent incident involving unauthorized access to shared storage revealed that a file transfer appliance installed within the Cabinet Office was being used. Cabinet Office, "Unauthorized access to shared file storage used by Cabinet Office staff" (https://www.cao.go.jp/others/csi/security/20210422notice.html, in Japanese).

*8 Antivirus technology these days generally updates patterns in real time to increase detection rates. This is also a big advantage because even if you can't detect a virus when the email is received, you may be able to detect it later when the recipient downloads or views the file.

These could be seen as weaknesses from an internal control perspective, however, so we can't leap on this option too enthusiastically. For example, many organizations commonly archive sent and received emails for tracing purposes. This provides important evidence in the event that an email exchange becomes a point of contention in a future litigation case or the like. But if attached files take the form of a URL, it may be difficult to trace exactly what files were sent and received.

Another risk is that it may create loopholes in measures to prevent information leaks and thus facilitate insider crime. Important documents could be uploaded to shared storage and taken outside of the organization via a one-time URL, with the file then deleted from shared storage. Administrators cannot realistically read through every email, so email auditing systems in some organizations archive emails for checking based on whether they have any attachments. But an email that only has a URL in the body text is no different from any other ordinary email, so it becomes very difficult to discover misconduct.

In the end, any system will have its advantages and disadvantages, so the key is how much of what sort of risks you are willing to accept.

Personally, I have been known to attach files to emails when I needn't have, so I can simply stop doing that, by which I mean reconsidering whether I really need to send certain information in an attached file, and if it really is necessary, sending it as an attachment and combining that with a temporary storage system and/or monitoring system depending on the risk. I think that is a reasonable middle ground. Either way, the call to abandon encrypted ZIP files will no doubt be an important turning point for the way we look at email systems and policies ahead.

1.4 Cautionary Notes on Online Conferencing Systems

As mentioned at the top, 2020 saw the rapid rise of telework. Online conferencing systems, only in use by a subset of users a year earlier, suddenly became widespread and reached a level of usability accommodating anyone and everyone. These online conferencing systems all share a particular mechanism in that users join a conference by clicking a one-time URL issued by the host. This system is very convenient because online conference participants simply need to click on the host's one-time URL received via email or the like.

But what if someone impersonates the online conference host and provides a phishing site URL instead? Users are accustomed to clicking on one-time URLs, so it's not hard to imagine them casually clicking on such a URL in an email. A fake login screen that looks identical to the genuine article could be used to steal user credentials including ID and password, allowing an attacker to retrieve internal company information and possibly gain access to other services if the user uses the same password for multiple services.

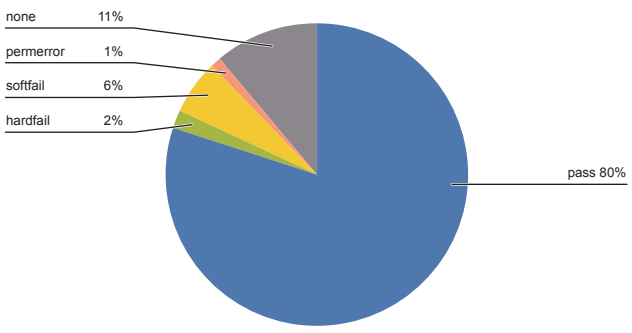

Figure 5: Breakdown of SPF Authentication Results

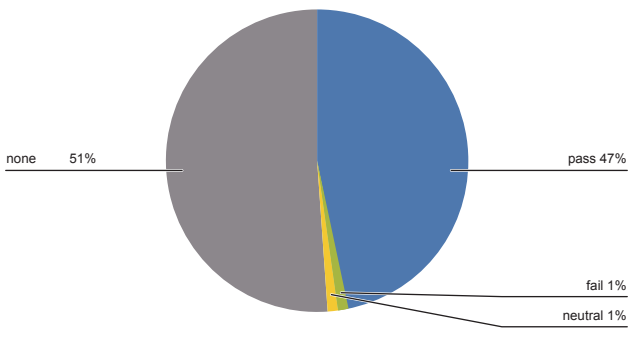

Figure 6: Breakdown of DKIM Authentication Results

Fortunately, email has a powerful framework called DMARC, which can authenticate the sender domain. Companies' and other organizations' domain administrators should ensure proper sender authentication is in place to combat spoofing.

Alongside system-based countermeasures, regularly reminding and educating users that some attacks start with emailed URLs can also further enhance the effectiveness of security measures.

1.5 The Rise of Sender Authentication

Figures 5 to 7 show breakdowns of sender authentication results aggregated across email services provided by IIJ over April 2020 – March 2021.

Sender authentication technology is undoubtedly effective against emails spoofing your domain and phishing emails impersonating well-known brands. The data here show sender authentication results as a proportion of all emails received, and it should be noted that large volumes of email, such as those sent out by mass email distributors, tend to swamp the rest of the data. Another point to note is that some spammers deal with sender authentication by buying domains for the purpose of sending out spam, so a "pass" result does not necessarily mean that an email is not spam.

Comparing these data with the results we reported in IIR Vol. 47 (https://www.iij.ad.jp/en/dev/iir/047.html), the proportion of "pass" results (successful authentication) is up a few percentage points for SPF, DKIM, and DMARC alike, and the proportion of "none" (no authentication information) is down for SPF and DMARC. DMARC has taken on the role of authenticating the domain found in the From header, so DKIM is really only used for authenticating email sent using an organization's own domain from a third-party SaaS, such as Salesforce and mass email distributors.

1.6 Conclusion

RFC822, specifying the current SMTP email protocol, was published in 1982. SMTP stands for Simple Mail Transfer Protocol, and it's hard not to be surprised that this same simple email protocol continues to run some 40 years later.

All sorts of communication tools such as online conferencing systems and text-based chat are popping up. But the reality is that email cannot replace all of them, so this situation is likely to persist for some time.

Looking ahead, IIJ will continue working to make the world safe for email.

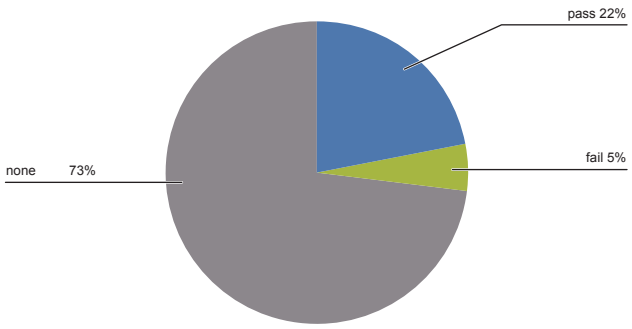

Figure 7: Breakdown of DMARC Authentication Results

Isamu Koga
Manager, Operation & Engineering Section, Application Service Department, Network Division, IIJ
Mr. Koga joined IIJ in 2007. He is engaged in the operation of email services and investigates email-related trends in the wild. To keep customers' email boxes safe, he communicates information about the latest attack methods, trends in spam, and countermeasures.

IIJ's Further Challenges as a Full MVNO

—Road to commercial NSA/SA services through a unique path in local 5G environments

2.1 Introduction

IIJ launched its MVNO[*1] services for enterprises in 2008 and for consumers in 2012. It has since offered a whole range of MVNO services and solutions as a leader in the industry and, since 2018, as Japan's first full MVNO[*2]. In terms of wireless networks, it initially used NTT Docomo's 3G network; in 2012 it became Japan's first MVNO to support 4G LTE; and in 2014 it added support for the KDDI 4G LTE network. It has thus evolved to become a multicarrier MVNO capable of providing optimal solutions to customers seeking carrier redundancy. And it has continued to grow using the latest wireless networks operated by MNOs[*3], adding support for KDDI's 5G network in 2020, for example.

So IIJ has always been at the forefront of the MVNO space, but its ultimate aim in rolling out a HSS (Home Subscriber Server)[*4] as a full MVNO is to advertise its PLMN (Public Land Mobile Network)[*5]—440-03 (IIJ)—from base stations and thus make it possible to offer IIJ's own services on an end-to-end basis including user devices.

To achieve this, IIJ needs to have not only a HSS/P-GW (Packet Data Network Gateway)[*6] but also an MME (Mobility Management Entity)[*7] / S-GW (Serving Gateway)[*8] and base stations. But telecommunications operators[*9] like IIJ have so far not really had any opportunities for frequency allocations that would allow them to obtain a commercial station license (wireless).

Against this backdrop, the report of the Information and Communications Council's New-generation Mobile Communications System Subcommittee (June 18, 2019) laid out the technical requirements for the 28.2–28.3GHz range from among the candidate frequency bands, and the necessary systems were put in place in December 2019. In terms of local 5G, the systems necessary for 4G communications systems that use the regional Broadband Wireless Access system (BWA)[*10] frequency band (2575–2595MHz) were also developed with the role of an anchor using NSA (non-standalone)[*11] architecture.

With the systems in place, it became possible for IIJ to own local 5G base stations / BWA base stations, and we quickly commenced the technical studies necessary for transitioning to NSA.

Section 2.2 below discusses our technical studies related to NSA deployment, Section 2.3 goes over some NSA deployment case studies, Section 2.4 looks at the functionality necessary for deploying an SA system, and Section 2.5 discusses our efforts to make the full VMNO concept a reality.

*1 Abbreviation of Mobile Virtual Network Operator.

*2 IIJ, "IIJ to Begin Offering 'IIJ Mobile Access Service Type I' as a Full MVNO" (https://www.iij.ad.jp/en/news/pressrelease/2018/0315-2.html).

*3 Abbreviation of Mobile Network Operator.

*4 A logical node that manages the subscriber information database for 3GPP mobile telecommunications networks. It manages authentication information and location information.

*5 In mobile telecommunication systems, a globally unique subscriber identification number called an IMSI (International Mobile Subscriber Identity) is issued to each device (internal SIM card). The IMSI is a combination of three identification numbers. The first three digits are the MCC (Mobile Country Number), which identifies the country/region, and the next two or three digits (depending on the country) are the MNC (Mobile Network Code), which identifies the operator. The remaining digits are the MSIN (Mobile Station Identification Number), which identifies the subscriber. The MCC and MNC combined are called the PLMN, which identifies the operator's network. MCCs are defined by the International Telecommunication Union (ITU-T), and Japan has been assigned two numbers: 440 and 441. MNCs are issued by the Ministry of Internal Affairs and Communications (MIC) based on applications from operators and consist of two digits in the range 00 to 99. The MIC has published guidelines for converting MNCs to three digits (https://www.soumu.go.jp/main_content/000663786.pdf, in Japanese).

*6 A logical gateway node and the transit link, the P-GW allocates device IP addresses, forwards S-GW packets, etc.

*7 A logical node that accommodates the LTE base station (eNodeB) and provides mobility control etc.

*8 A logical gateway node that accommodates the 3GPP access system.

*9 A telecommunications carrier that does not own transmission line equipment (previously known as a Type 2 Telecommunication Carrier).

*10 Wireless systems for telecommunications services that use the 2.5GHz band, created with the objective of improving the public welfare across regions by, for example, improving public services and closing the digital divide (gap faced by areas disadvantaged by poorer access to modern ICT).

*11 Architecture defined in 3GPP Release 15; the industry mainstream is Option 3x, whereby LTE is used for control signalling and 5G is used to send data signals. For the core system, the LTE MME/S-GW/P-GW are used.

2.2 Technical Studies for the Shift to NSA (Non-Standalone)

The equipment required for transitioning to NSA is: HSS, MME, S-GW, P-GW, 4G anchor (BWA) base stations, and local 5G base stations. Having become a full MVNO, IIJ already has its own HSS, so we have already done the HSS development work needed for transitioning to NSA. We have also developed the HSS functionality shown in Figure 1 to prevent users on BWA-only contracts from using 5G.

The provisioning interface between the BSS (Business Support System)*12 and HSS also needs to be changed, but we were able to deal with this via a minor expansion to the existing full MVNO interface. As an MVNO with LTE support, we already had a P-GW. We had the option of only setting up a new MME/S-GW, but sharing the P-GW among existing IIJ mobile services and NSA services would complicate the P-GW resource design and make it difficult to isolate when faults occur, so we decided to set up the MME/S-GW/P-GW anew.

We also looked at whether it would be possible to make use of IIJ's existing equipment for the Radius / PCRF (Policy and Charging Rules Function)*13 / OCS (Online Charging System)*14 / OFCS (Offline Charging System)*15, which control the S-GW/P-GW. In light of the impact on IIJ's mobile services, we decided to develop a new system, as with the P-GW. We used a little ingenuity in designing the Radius. Giving the PCC (Policy and Charging Control) rules managed by the PCRF to the Radius means the PCRF is no longer needed. Since the PCRF is no longer required, neither is the OCS for online processing. So we only needed to develop the Radius/OFCS, which greatly reduced the cost and number of steps. But there is a disadvantage with this. Because we can't do real-time processing, speed limits and the like are processed in batches, so there is a one-day lag when imposing them.

In terms of services, we are now able to provide unlimited-capacity services in both BWA and 5G. Unlike IIJ mobile services, the services do not use carrier equipment, so in effect they move us toward the ultimate aim for IIJ as an MVNO, which, as mentioned in Section 2.1, is to make it possible to provide IIJ's own services on an end-to-end basis. We also developed dual connectivity*16 (downlink only) for BWA and local 5G and thus provide services that aggregate BWA speeds and local 5G speeds.

We also need to lay the backhaul line from the base stations to the MME/S-GW. Mixing the C-Plane, U-Plane, and M-Plane*17 all into a single line can result in the crucial

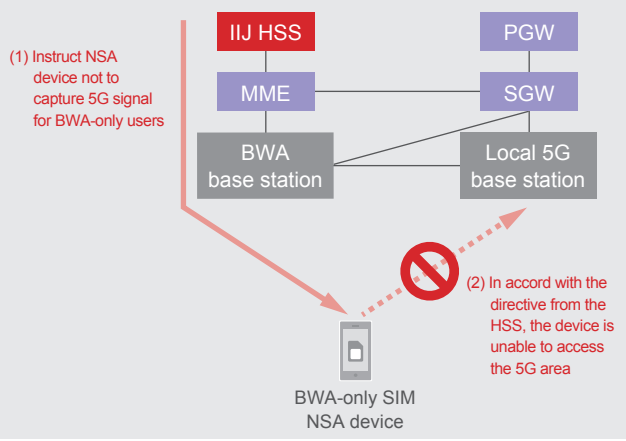

Figure 1: Mechanism for Preventing BWA Customers from Using 5G Services

*12 General term for business support systems for telecommunications carriers. Manages customer information and charging (billing) information, and supports business processes such as applications for services, commencement of service, billing, and inquiries.
*13 Logical node that controls QoS and billing for user data transfers.
*14 Online Charging System.
*15 Offline Charging System.
*16 A technology for improving speeds on wireless network sections by bundling multiple component carriers between separate base stations.
*17 Control Plane (C-Plane) / User Plane (U-Plane) / Management Plane (M-Plane).

C-Plane packets being discarded, so we decided to separate them with a VLAN and assign priority based on the DSCP (Differentiated Services Code Point) value.

As for the SIMs (Subscriber Identity Modules)[*18], we arranged two types: one that can only use the BWA base station area, and one that can use the Docomo LTE network area and the BWA base station area. Japan's Guidelines for Local 5G Implementation say that while it is not possible to link local 5G networks to complement nationwide MNO services, it is possible to use nationwide MNO services for the purpose of supplementing local 5G services. While this can be interpreted to mean that it is possible to link together local 5G areas and nationwide MNO 5G areas, we determined that this would necessitate coordinating with the MNOs, so we only arranged SIMs capable of using local 5G areas.

For BWA base stations, we developed an adaptive modulation method that switches among QPSK, 16QAM, 64QAM, and 256QAM[*19] depending on radio quality, with our requirement being to achieve radio speeds above those of BWA base stations provided by other operators. Developing this made it possible to achieve up to 295Mbps (256QAM) wireless downlink throughput and up to 17Mbps (64QAM) wireless uplink throughput under MU-MIMO (Multi-User MIMO)[*20] 4x4 conditions. As for local 5G base stations, only 100MHz of bandwidth in the 2GHz band was available for licensing, but we were able to achieve up to 484Mbps (64QAM) wireless downlink throughput and up to 125Mbps (64QAM) wireless uplink throughput under MU-MIMO 2x2 conditions. Based on the above technical investigations, we were able to complete IIJ's NSA architecture (Figure 2).

2.3 NSA Implementation Case Studies

(1) Shiroi Wireless Campus

IIJ's mobile technology testbed, Shiroi Wireless Campus in Shiroi-shi, Chiba Prefecture, went into full operation in November 2020. More than just a showcase for the latest mobile technology, Shiroi Wireless Campus will also be made available as an environment for testing interoperability between user devices and network equipment as well as between different pieces of network equipment, and as an environment for proof-of-concept testing in collaboration with customers in the aim of making use of the latest mobile technology.

A highlight of Shiroi Wireless Campus is local 5G, for which a commercial station license was obtained in March 2021. As explained in Section 2.2, we have already built the NSA core housing the local 5G base station / BWA base station based on NSA architecture design concepts for IIJ. We

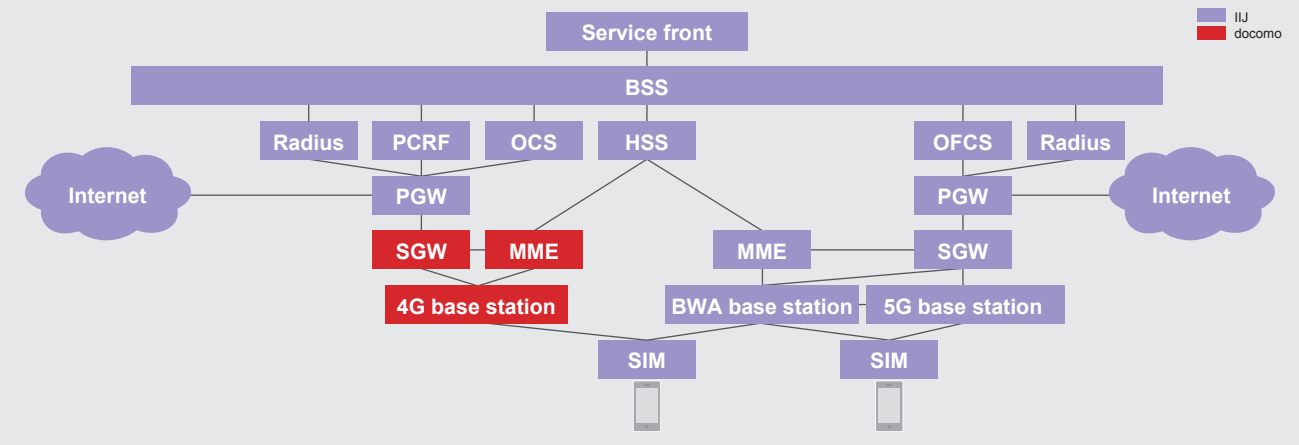

Figure 2: IIJ's NSA Architecture

*18 Officially called a UIM (User Identity Module Card) or USIM (Universal Subscriber Identity Module Card) but generally referred to as a SIM card.

*19 Wireless modulation methods.

*20 A method of increasing the amount of data that can be transmitted simultaneously using spatial multiplexing technology, whereby multiple antennas transmit data at the same time and multiple antennas receive the data. With MU-MIMO, multiple users can transmit at the same time. In contrast, only a single user can transmit with SU-MIMO (single-user MIMO).

developed the base stations / MME so that each of the base stations could advertise multiple PLMNs (see Figure 3). The local 5G base station and BWA base station are now in operation and emitting radio waves (see Figure 4).

Alongside local 5G / BWA, we are also building a heterogeneous wireless environment at Shiroi Wireless Campus. We have already connected the sXGP (band41) base station to the MME/S-GW, and we also intend to build Passpoint*21-compatible Wi-Fi access points into the HSS. To facilitate seamless communications between different wireless areas, we incorporated SIM design knowhow gained through proof-of-concept work*22 with the University of Tokyo's NakaoLab into IIJ SIMs. As for the SA (standalone)*23 system, we plan to obtain a Sub-6*24 base station / SA core and build them at Shiroi Wireless Campus.

■ (2) Grape One

To launch local 5G services, we established a new company, Grape One Co., Ltd.*25, in association with Sumitomo Corporation and CATV companies. The knowhow gleaned from the NSA technical studies described in Section 2.2 is also being used in Grape One's NSA core.

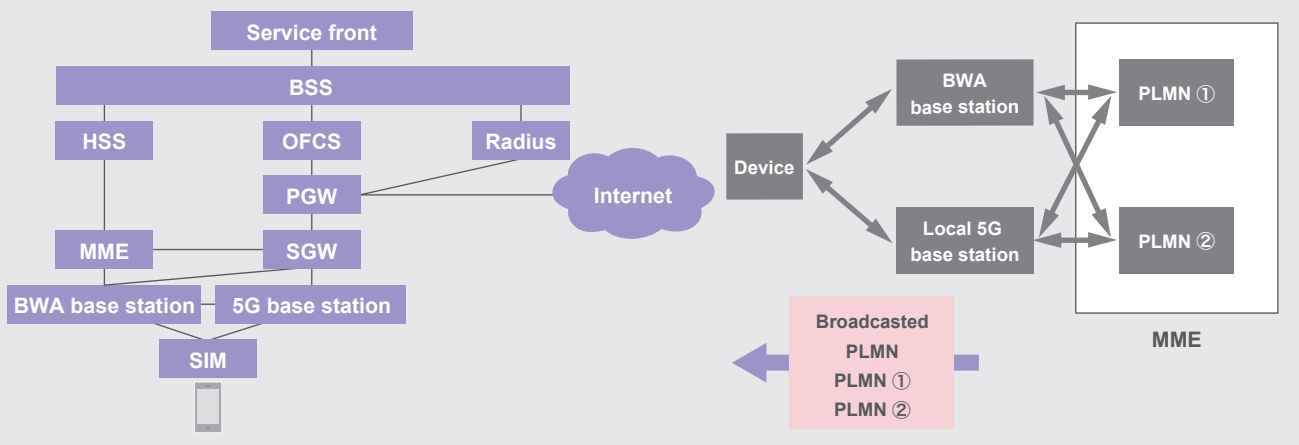

Figure 3: Overview of NSA Network Structure and Multiple-PLMN at Shiroi Wireless Campus

Figure 4: Local 5G Base Station and Device Speed Test at Shiroi Wireless Campus

*21 Communication standard certified by the Wi-Fi Alliance, which manages wireless LAN standards. Formerly called HotSpot 2.0. Normal Wi-Fi requires an SSID and password, but with support for this standard enabled, a device will automatically connect upon entering a Wi-Fi spot.

*22 IIJ, "First in Japan! University of Tokyo and IIJ start proof-of-concept tests on integration of public LTE and private LTE" (https://www.iij.ad.jp/news/pressrelease/2019/0605.html, in Japanese).

*23 Architecture in which 5G base stations operate on a dedicated 5G core network independent of the LTE core network (HSS/MME/S-GW/P-GW). 5G SA Option 2 is the industry mainstream option.

*24 Frequency bands below 6GHz.

*25 IIJ, "Pursuing the Wireless Platform Business Using Local 5G" (https://www.iij.ad.jp/en/news/pressrelease/2019/1224.html).

2.4 Functionality Necessary for Implementing SA (Standalone)

As mentioned in Section 2.3, we have already started commercial services with the NSA core and local 5G base stations / BWA base stations. The next necessary step is to look at how to incorporate SA into IIJ services. We therefore investigated what functionality would be needed to implement SA.

(1) NEF (Network Exposure Function)[*26] to control SA core from applications

3GPP R16 provides quite a bit more clarity than R15 on the functionality required of the NEF (see Figure 5). 3GPP TS 23.502 contains an updated list of NEF APIs, and while it is based on the LTE SCEF (Service Capability Exposure Function)[*27], it also features APIs not present in SCEF, such as Nnef_TrafficInfluence/Nnef_ AnalyticsExposure. IIJ has a bevy of cloud applications, such as an IoT platform called the IIJ IoT Service, so it will be possible to provide some interesting SA services using the NEF.

(2) Centralization of subscriber profiles provides single point for queries

With NSA, subscriber profile data was distributed across different equipment (UDR/MME/S-GW/P-GW), so subscriber profiles had to be queried at multiple points and processing this data was complicated. So 3GPP TS 23.501 "4.2.5 Data storage architectures" defines the Nudr interface for centralized storage and retrieval of previously distributed data on a UDR (Unified Data Repository) / UDSF (Unstructured Data Storage Function) (see Figure 6). This enables cloud-based centralized data management and provides a unified point (UDR/UDSF) for querying subscriber profiles, making data processing more efficient.

(3) Local breakout (MEC[*28])

NSA does not have a local breakout function, so the S-GW/P-GW had to be distributed around the area. So with SA, we physically separated the logical functions within the S-GW/P-GW. That is, we newly defined and physically divided the SMF (Session Management Function), which controls the session function, the PCF (Policy Control function), which controls the PCC rule function, and the UPF (User Plane Function), which controls packet processing.

Consolidating the SMF in the cloud means it has a major role in maintaining sessions from old UPFs to new UPFs, making breakout easier. If the SMFs were distributed, session information data would have to be synchronized across each of them, so there would have been no point in physically separating the system components. Consolidating the PCF in the cloud is a significant move for the same reason as with the SMF.

Also, using the NEF API Nnef_TrafficInfluence, as mentioned in (1) above, enables local breakout via application-based control.

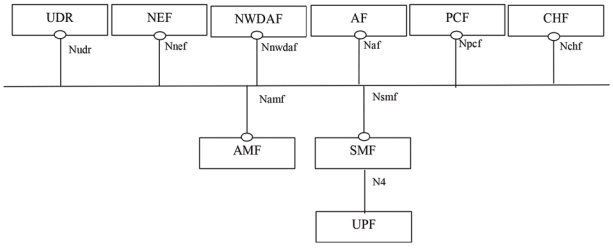

Figure 5: SA Core Architecture Including NEF (excerpt from 3GPP TS23.503)

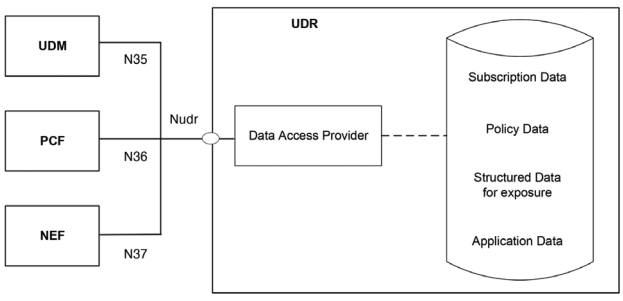

Figure 6: Data Storage Architectures (excerpt from 3GPP TS 23.501)

*26 Logical node that provides an API (application programming interface) to external applications for controlling the logical network nodes (Network Functions: NFs) that make up the 5GC (5th Generation Core network). In addition to being able to see subscriber information and network state changes in detail on the application side, it is also possible to control NFs via applications.

*27 3GPP TS 29.122-compliant logical node that provides an API to external applications for controlling LTE core network components such as HSS/MME.

*28 Abbreviation of Multi-access Edge Computing.

(4) Network slicing

NSA already has technology for network slicing between devices and the core network with multiple PDP Contexts[*29] and dedicated bearers[*30], but with SA it becomes possible to specify detailed SLAs (Service Level Agreements) for each network slice. While RAN (Radio Access Network) network slicing is still under development, support on the RAN side is essential in order to enable network slicing between devices and the UPF.

(5) 5G-AKA

LTE had incomplete IMSI encryption, so 5G-AKA introduces a mechanism for encrypting the MSIN (Mobile Station Identification Number) of the 5G subscriber ID (SUPL) that is equivalent to IMSI. IIJ has also succeeded in getting 5G-AKA-compatible eSIMs to work on the SA core[*31]. For details on 5G-AKA, see 3GPP TS 33.501 V17.1.0 (2021-03).

Looking ahead, we plan to procure Sub-6 / SA core products that meet the functional requirements described in (1) to (5) above and build the systems at Shiroi Wireless Campus.

2.5 The Path to Full VMNOs

With nationwide 5G areas, MVNOs will no doubt have to connect to nationwide MNO base stations, as is the case with full MVNOs, because only MNOs can obtain frequency licenses.

IIJ is advocating the concept of the full VMNO[*32] and envisions RAN sharing[*33] for nationwide MNO 5G base stations. With RAN sharing, operator PLMNs are distributed as keys, so MVNOs also need to have an SA core.

Local 5G bases stations, meanwhile, are owned by the individual operators, so there is a wider range of options, including setting up your own SA core or connecting to a VMNO's SA core.

As explained in Section 2.4, we plan to procure an SA core and Sub-6 base stations for Shiroi Wireless Campus so that we can test them and identify any issues as part of our efforts to make VMNOs a reality.

2.6 Conclusion

COVID-19 has accelerated moves toward remote work and other DX (digital transformation) initiatives. Before the pandemic, people expected the rollout of local 5G to happen after the Tokyo Olympics and to take several years, but it now looks to be moving ahead earlier than that.

Having installed BWA base stations, local 5G base stations, and an NSA core, IIJ is now also in a position to develop its own wireless and core system knowhow.

With an eye to rolling out an SA system ahead, we will continue to coordinate our efforts not only with mobile services but IIJ's other services as well so that we can continue to provide the types of services not available from other providers.

Jun Kakishima
Technology Development Department, MVNO, IIJ.
Since joining IIJ in 2017, Mr. Kakishima has worked on the development of corporate mobile services. When IIJ launched its full MVNO operations, he worked on specifications development, construction, and operations design for the core equipment, including the HSS. He is also involved in specifications development, construction, and operations design for local 5G services and NSA/SA, and in planning the introduction of SoftSIM/eSIMs.

*29 Providing multiple PDP (Packet Domain Protocol) contexts to a single user (device). See 3GPP TS 23.976.
*30 Setting up multiple bearers for a single PDP (Packet Domain Protocol). See 3GPP TS 23.401.
*31 IIJ, "Establishing the essential technologies needed for full MVNO and local 5G services, beyond the 5G core networks technologies" (https://www.iij.ad.jp/en/news/pressrelease/2020/1102.html).
*32 Internet Infrastructure Review (IIR) Vol.48, Focused Research (1), "MVNOs in the 5G Era: Advocating the VMNO Concept" (https://www.iij.ad.jp/en/dev/iir/048.html).
*33 Multiple operators' core networks sharing a single radio access network (RAN).

Meet Barry, IIJ's Tool for Rapid Fault Resolution

3.1 Background to Barry's Deployment

To provide stable, high-quality services, IIJ must attend to a range of operational tasks. Key among them is troubleshooting, which involves restoring service-providing systems when they are unable to maintain a normal operating state because of hardware or software faults. IIJ uses an internally developed operations system called Barry for troubleshooting. Here, I describe how Barry works and what it does.

First, however, I would like to talk about how we dealt with faults before Barry. Normally, the services we provide are constantly monitored for any anomalies in the equipment and functionality provided. When service anomalies arise, an alert is generated, prompting us to take action. The first step is to find troubleshooters capable of dealing with the issue. At IIJ, we call this escalation. Next, the troubleshooters begin the job of restoring the service to a normal state. The actions taken differ depending on the service, but generally the process involves the people involved communicating with each other and using various tools to investigate and record the issue as they work toward a solution.

This is how IIJ had been dealing with faults until now, but this approach had its problems. So we revised our approach and developed an operations system called Barry to facilitate smoother troubleshooting.

The two main problems with the previous approach were as follows.

■ Problem 1: Finding troubleshooters and accurately communicating the details of the issue

We need to find people quickly. We currently still use phone calls to contact candidate troubleshooters and ask if they can tackle the issue. The advantage of using the phone is that we can call them continuously. Emails and other messaging can also be used to escalate an issue, but such messages are usually only sent to candidates when the issue arises. With the phone, however, we can continue calling the person until we reach them, so we have a higher success rate in getting hold of the right people for the job. The flip side of this is that the people manually making the phone calls are tied up until the troubleshooters are found. This issue can be addressed by using automated phone calls, but this entails one-sided voice communication, so it can be difficult for the person to clarify and confirm the details, and the issue of automated calling system cost also remains. And there are limits on how many people can be called at once.

Also with phone calls, people can mishear or fail to hear what was said, and communicating English abbreviations and symbols is also difficult. An advantage with email-based escalation, however, is that these issues do not arise and it is easy to communicate complicated information.

Based on the above, we identified the ability to call people continuously and to accurately communicate information as two key points. Solving these issues should speed up the initial part of the troubleshooting process.

■ Problem 2: Reducing the load on troubleshooters

The work of dealing with faults puts a load on the troubleshooters in several ways. Systems can recover from some simple faults automatically, but the faults we are talking about here are those that require a troubleshooter with deep knowledge of the service to tailor a response to the situation. So high-level knowledge of the service and the right skills to address the issue are required. Other sources of pressure on troubleshooters include the need to sometimes respond on holidays or at night and demands for the rapid restoration of service. And on top of the high difficulty of dealing with faults, they need to communicate and share information with other concerned parties.

Under these circumstances, it was often the case that certain individuals, depending on the system, would handle much of the work. But it was difficult to ascertain how much of a skew there was in the workload. To enable troubleshooters to concentrate on dealing with the issue, we wanted to make it as easy as possible for them to perform the peripheral tasks.

3.2 Addressing the Problems

We thus looked at using an operations system to enable a fast response and reduce the load on troubleshooters. Adopting an existing tool was also an option, but none of the tools available were immediately suitable to IIJ's own response process, so with an eye to optimizing the internal workflow, we decided to develop a system in-house. While this approach requires cost outlays, it also has a strong advantage in that we can continuously improve the system as needed.

Firstly, we had the idea of building it as a smartphone app to solve the problem of getting hold of people. The concept was to mimic the incoming call screen and to display a text message once the "call" had been answered, thus realizing the advantages of both modes of communication. This simultaneously has the advantage of messaging, in that it is text-based, and the advantage of phone calls, in that people can be called continuously. It also does away with the limitations of phone calls , opening the door to the notion that we could customize the system in terms of the escalation sequence, how many people are called at the same time, and so on. We saw the potential flexibility to tailor the call to the style of the operations team.

In terms of reducing workloads, we felt we needed to be aware of what aspects of the process troubleshooters find inconvenient, so we told several operations personnel about the idea of using smartphones to page them and asked for their thoughts. Many of the responses indicated that sharing information was troublesome. The process of dealing with faults involves understanding the details of the fault, sharing information when implementing the response, and incident tracking. At the time we spoke to the operations personnel, each operations team was using different tools to share information in real time when implementing a response, including IRC (Internet Relay Chat) and SaaS communication tools. Various methods were also used to record the faults as incidents, including email and Wiki/ticket systems. Another inconvenience was knowing that a fault has occurred but being unable to tell what the response status is when away from the PC screen. Given these points, it was apparent that we needed to integrate information sharing and tracking for the people to whom issues are escalated. This is because tool ease of use has a major impact on the efficiency of the response process. Our focus with the new operations system was on usability, with the opinions of troubleshooters taken into account.

3.3 Barry's Featuresn

We started implementing the new operations system under the name Barry. The name comes from the famous Swiss mountain rescue dog and signifies our hope that the tool will come to the aid of those dealing with system faults.

Barry's features are divided into three parts: server, Web frontend, and mobile app. The server implements core functionality, such as escalation and incident tracking, and exposes it as a gRPC API. The Web frontend and mobile app use the server's API to provide a UI (Figures 1

Figure 1: Barry's Mobile App Screen

& 2). Conducting the entire troubleshooting process via the mobile app would be a bit daunting at present, so we have each part doing what it does best. We see the mobile app as mainly being for calling people and facilitating simple communication, and we have structured the system on the premise that the bulk of the incident response will happen on PC via the Web frontend.

Using smartphones as a tool makes it possible for people to offer advice and other support in circumstances when previously they would not have been able to tell what was happening or be involved in the response. The mobile app is made available internally through a mechanism for distributing apps within an organization.

I will now go through specific features we implemented in Barry.

Feature 1: Flexible calling

To implement the smartphone calling feature, we used the same technology as an ordinary phone call app. When the server is given a request to initiate escalation, it sends a smartphone notification to the operations team for the service on which the fault has occurred (Figure 3). Upon receiving notification that an escalation has been initiated, the mobile app displays the incoming phone call UI. Once users answer the call, they launch the mobile app and review the details recorded on the server; they then reply via the app to say whether they can deal with the issue or not. Once an affirmative answer is received, the escalation process is complete. This is the basic mechanism, and the operations team can freely configure the server for the desired call order, number of devices called simultaneously, ring time, and number of retries.

There are also two patterns for initiating escalation: automatic and manual. An escalation can be generated automatically in response to a service monitoring alert, which I mentioned above. The system also supports manual escalation so people can be called in emergencies independent of whether an alarm is generated.

Feature 2: Integrated information tracking

An issue we identified was that dealing with faults was burdensome for troubleshooters because they were using all sorts of tools in the process. Barry provides functionality that integrates the entire process from escalation to incident tracking. In addition to the calling feature, we also implemented an incident tracking mechanism. The tool is functionally equivalent to an issue tracking system and allows people to record details of the fault and track response status.

Specifically, each failure is deemed to be an incident, and the operations team's communications with each other and updates to response status are recorded up until the issue

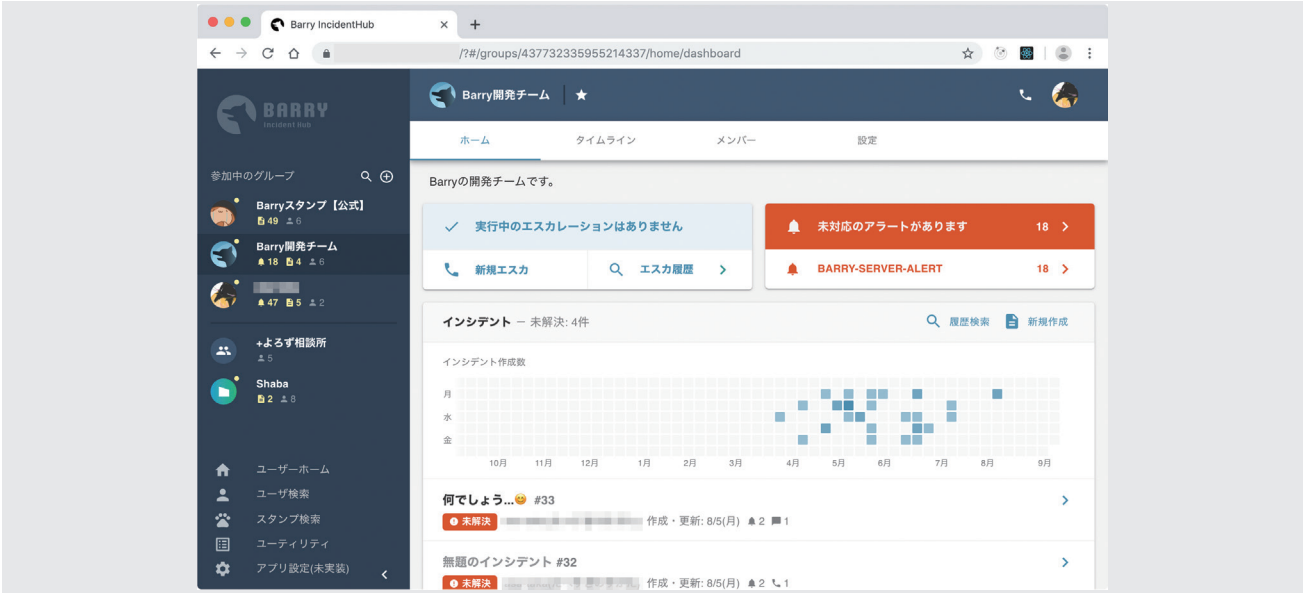

Figure 2: Barry's Web Frontend

is resolved. Alerts generated can be linked to incidents. This recording of incident histories makes it possible to refer to past examples when addressing faults.

Implementing this functionality made it possible to handle the entire process from fault occurrence through to resolution via a single tool. And because Barry supports both Web and mobile app interfaces, frontline troubleshooters can view the status and make comments even while on the move.

Ease-of-use considerations

The benefits of implementing functionality to integrate a range of tools would be limited if it ultimately resulted in lower efficiency. We therefore made ease of use a priority with Barry's tools.

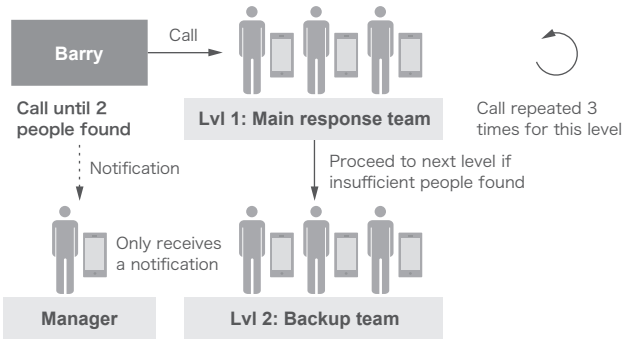

Figure 3: Barry's Calling Process

We interviewed troubleshooters when designing the system, and we then created mockups and asked for feedback to ensure we were on the same page. Repeating this process several times clarified what features were needed, and it also gave us early feedback on usability. We created a lot of fine-grained functionality, so here I will run through the major features.

Activity history display

To make it easy to see how often a phenomenon occurs and how much work is involved in rectifying the fault, we implemented statistics and visualization. This feature graphs a time series of alerts and activity for each user (Figure 4). Displaying alerts on a timeline enables efficient analysis of the circumstances under which faults are occurring. And making it easy to see the operations team's activity history helps managers understand what is going on more accurately than before.

The timeline display feature shows events in order of occurrence. When using Barry, users see a lot of alerts and new incidents/comments. It's not uncommon for users to have multiple operations teams, and it can be difficult to understand what is happening when many events occur at once. The timeline feature displays events for each user in chronological order, making it easy to keep track.

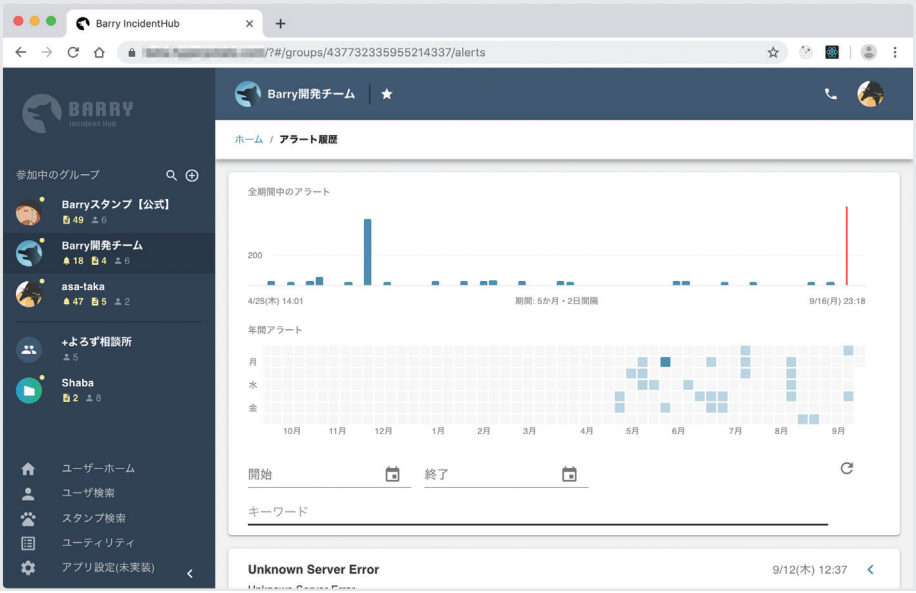

Figure 4: Graph of Alert Frequency

Stickers

We created emoji reactions and stickers for the incident comments to streamline communication (Figure 5). One problem is that expressions of gratitude and other emotional content in the form of comments increase the amount of information relating to an incident, making it hard to pick out the important details. With this in mind, we implemented an emoji reactions feature. We also created stickers like those used on social media, making it easy to convey basic/standard informational content.

Avatars

Avatars can be configured for each user and operations team. This feature is also widely used on social media services and helps improve visibility. The purpose is to prevent mistakes by allowing users and operations teams to freely configure their own avatars.

Webhook

The available features are also designed with automation in mind. An API is provided for everything that can be done via the screen, so users have the option of automating via software. Barry also has other automation features, notably webhook, which we implemented in response to user requests. Webhooks are a way for Web applications to provide information to external systems and are widely used by Web services and the like. Barry acts as the recipient of this information and thus supports the receipt of alerts and escalation initiations. Specifically, linking to webhooks such as Grafana makes it possible to link into existing systems without additional development. We also created a command line tool, so Barry can be used via simple scripts. We expect these features to be used in automating the work involved in dealing with faults.

3.4 Using Barry to Deal with Faults

Now let's follow the system operations process with Barry deployed.

Barry is an operations system for use within IIJ, and service operators perform the following initial setup.

Figure 5: Example of Emoji and Stickers on the Comment Screen

1. Set Barry to be the destination for service monitoring alerts
2. Define rules to say who is called and in what order when an issue is escalated
3. Troubleshooters install the Barry mobile app on their smartphones

With these arrangements in place, service monitoring alerts are sent to Barry when they occur. Upon receiving an alert, Barry saves the details, determines which operations team to use, and escalates the issue. The escalation process involves ringing people's smartphones according to the calling rules defined for the chosen operations team(s) until a troubleshooter is found.

Users learn of the escalation when their smartphones ring and then check why it was raised. The notification includes details of the alert, and if they are able to deal with the issue, users reply via the app to say they will start working on it. The system ends the calling process at this point, and the group is notified that responders have been selected.

The escalation feature is done with its role at this point, and the focus shifts to the incident features that provide integrated information tracking. The troubleshooters go over the event based on the information in the alert and put this information together into an incident. They then start working on rectifying the fault, leaving comments as they go. Information is shared within the operations team as it is added, including notifications to the mobile app, and people other than the designated troubleshooters can also add comments as necessary. Additional people can also be called on if the troubleshooters are unable to handle the issue alone.

Once the fault has been dealt with, the incident is updated as complete and Barry's work is done. The information recorded on the incident and the escalation history are stored in the system. A search feature is also available, so responders can refer to how similar issues were handled in the past as they work to fix a fault.

3.5 Operations

Barry's system is needed when dealing with faults in a range of services, so it needs to have high availability. Naturally, Barry itself can also fail, so it is designed and operated on the assumption that faults will occur.

The system is structured to have three independent regions, two of which provide redundancy for a single Barry system (Figure 6). The remaining region runs a separate Barry system. This is used by Barry's operators and comes into play when dealing with faults in Barry itself.

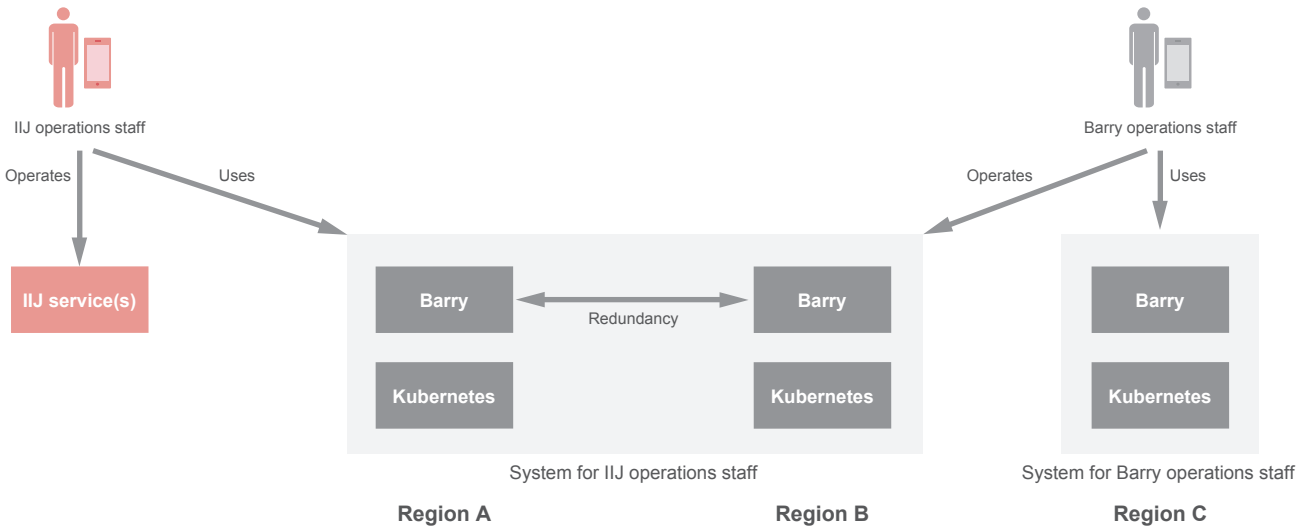

Figure 6: Barry System Structure

An independent Kubernetes cluster is run in each region, and Barry runs atop Kubernetes. Configuring the system to use Kubernetes' features eliminates the need to worry about hardware failures when running Barry.

Barry uses an external service to send notifications to smartphones. It is implemented as a combination of multiple external services so that a failure or delay in an external service does not become a single point of failure (Figure 7). The server looks at the responses of smartphones to which notifications are sent, and if an anomaly in the notification system is detected, Barry automatically falls back to using automated phone calls.

In the event of a top-level domain failure, the service may become inaccessible due to a name resolution failure, even if Barry is operating normally. To address this, we have set up multiple domains to ensure service access redundancy.

While a little different from system faults, we also deal with mobile app problems. On the server side, operators can roll back when problems occur, but they are not able to deal with issues in the apps installed on individual devices. Fatal errors cause the calling functionality to stop working, so two versions of the app are distributed. Along with a normal version of the app that is updated from time to time, an emergency version that is confirmed to be stable can also be installed.

3.6 Deployment and Impact

We released Barry internally in July 2020. Replacing the entire fault response system all at once would not be realistic, so our approach since the release has been to switch individual services over to Barry for the teams that want it. On the user end, there was the need to replace the mechanism by which alerts are sent and so forth, and this work of switching things over is progressing with help from the service teams. It's quite easy to adopt Barry particularly for newly launching services.

There are also tools created by Barry users now, so we have a real sense that the decision to open up the API is helping to facilitate the automation and streamlining of work for users.

Barry enables continuous calling like phone calls while also efficiently communicating information by sending a text message at the same time. Automation makes it possible to leave the simple procedure of calling people up to Barry. And the operating structure increases parallelism in the calling process, so candidates can be contacted all at once. When phoning people one after the other in sequence, we

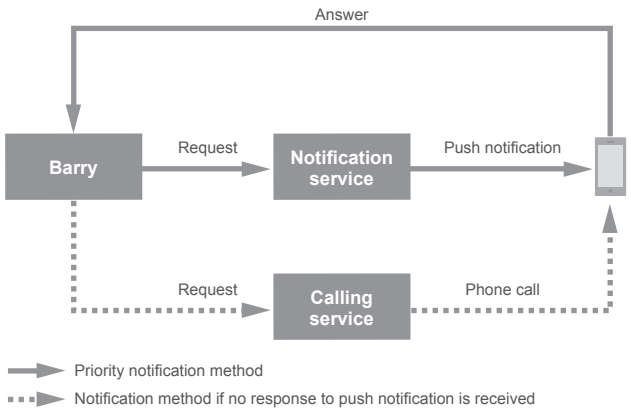

Figure 7: Fallback for Barry System Faults

experienced delays in initiating the response when only some of the candidates were able to deal with the issue, and we have been able to resolve this as well. One of our tasks was to speed up the initial part of the troubleshooting process, so the system helps with this.

There are currently 633 users and 190 operations teams. We conducted a post-release user questionnaire about using Barry. We asked whether frequency of use and Barry's introduction had improved the way they work, and we also asked what users would like to see improved. I will go over these topics below.

As to what had improved, the responses mentioned the speeding up of the initial response and the ability to see current status, which were tasks we had identified. Having the system make the calls has reduced the workload, and automating the process from alert to escalation means that people can find out about faults happening earlier. The responses also mentioned communication. Because it is now easier to tell what the status of the fault response is within operations teams, people are finding it easier to coordinate their efforts. The positive feedback on improvements flowing from Barry's introduction indicates to us that it is lending a hand on the operations front.

Meanwhile, some people have also asked for improvements to Barry.

One request is to simplify links with existing systems. We have provided an API and designed Barry to be suitable for a range of use cases, but modifications do need to be made to existing systems in order to use Barry. Users have asked us for a way to get started using Barry with only minimal changes to existing operations systems. The system is designed to work within IIJ's own unique set of circumstances, so we plan to address such individual requests in a flexible manner going forward.

Another was to address concerns about stability. As discussed, Barry is a system that is used when faults occur, so it needs to be stable. One criterion users look at when assessing a system for adoption is its track record in operation, but having been released not long ago, Barry lacks an adequate track record. To ensure people can use the system with peace of mind, a priority for us in providing Barry is to build up this sort of stable track record ahead.

Deploying Barry internally was a major milestone for us, but we still have work to do. We hope to contribute to maintaining and enhancing the quality of IIJ services by continuing to update the system going forward.

Yushi Nakai
Operation System Development Section, Operation Engineering Department, Infrastructure Engineering Division, IIJ
Mr. Nakai joined IIJ in 2007. He is involved in the development of services and operations systems.

July 2021 Vol.51

Internet Initiative Japan

About Internet Initiative Japan Inc. (IIJ)

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

IIJ-MKTG020-0049

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/