IIR

Internet
Infrastructure
Review

Mar.2018
Vol. 38

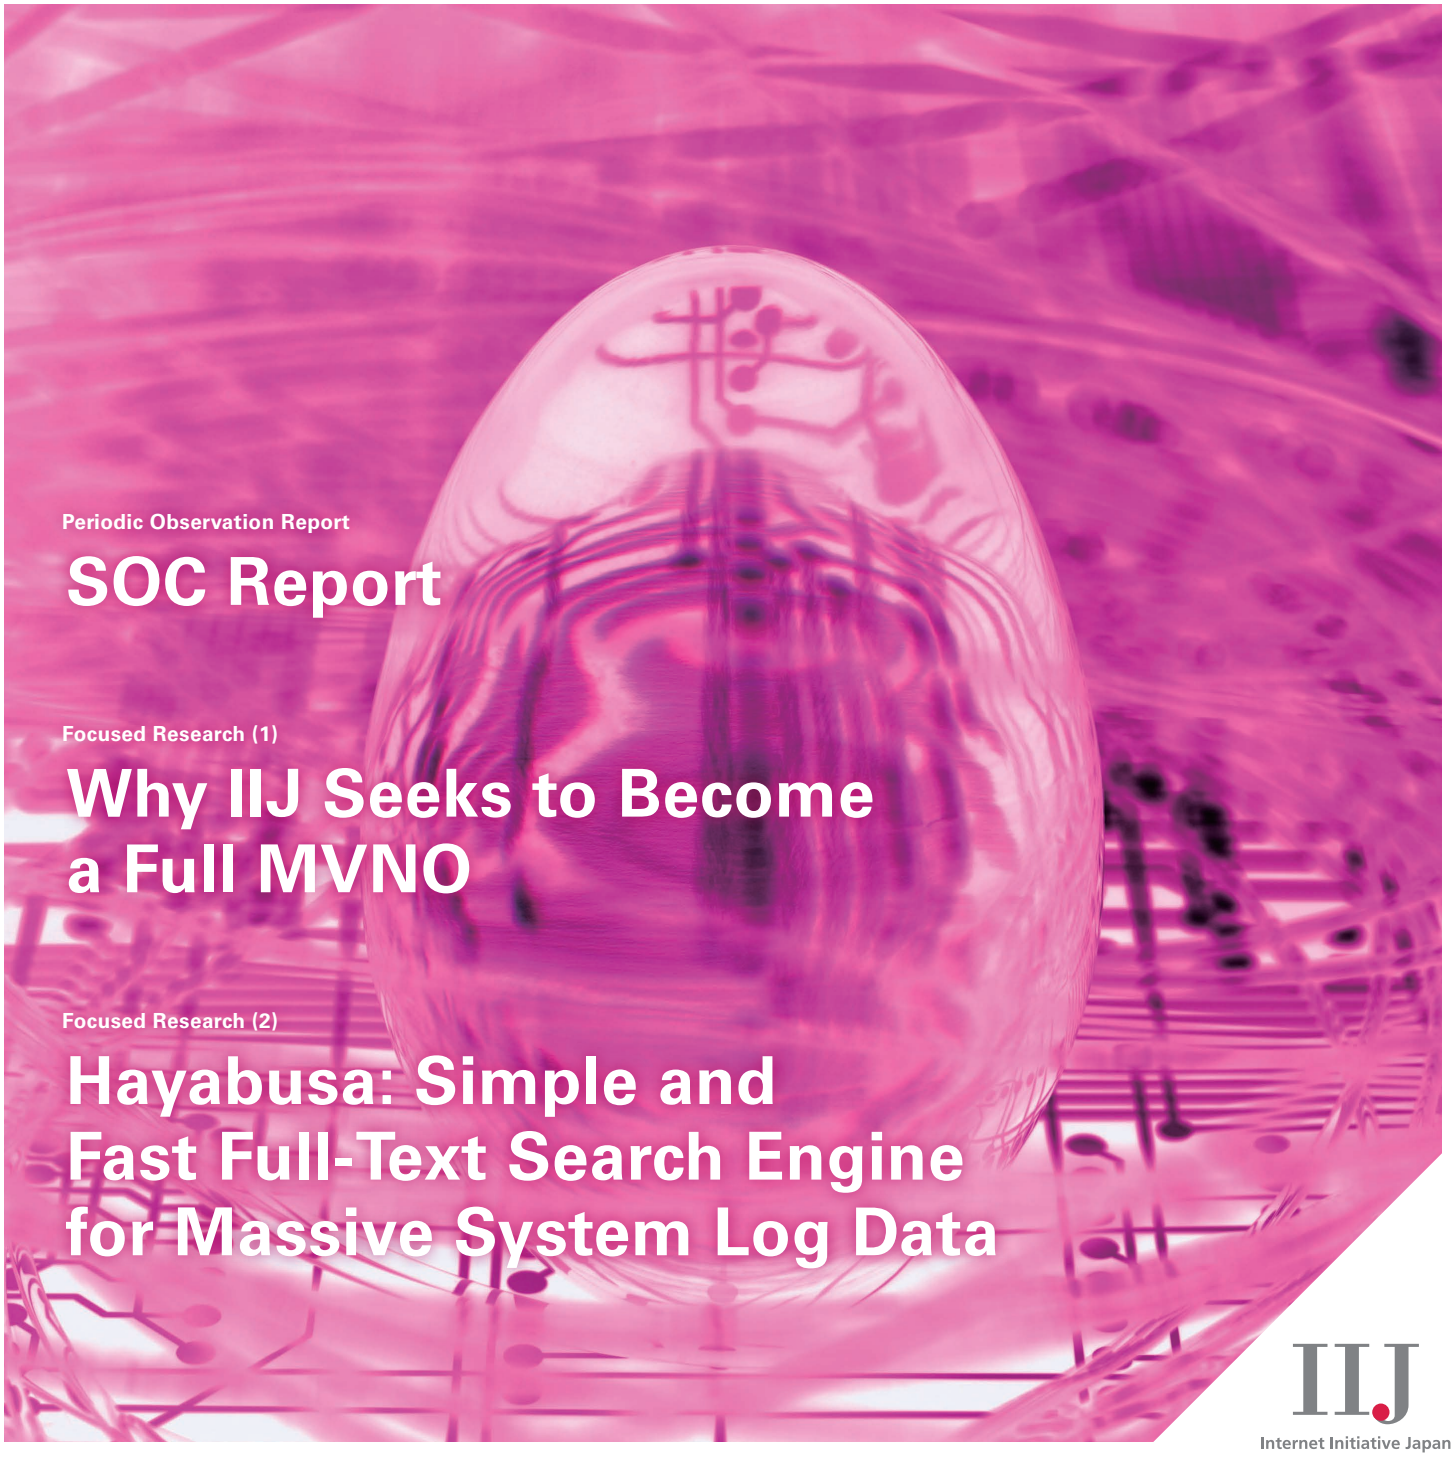

Periodic Observation Report

SOC Report

Focused Research (1)

Why IIJ Seeks to Become a Full MVNO

Focused Research (2)

Hayabusa: Simple and Fast Full-Text Search Engine for Massive System Log Data

IIJ
Internet Initiative Japan

Internet Infrastructure Review

March 2018 Vol.38

Executive Summary

Over the past three months, there have been a series of reports on cryptocurrencies being misappropriated from cryptocurrency exchanges, which has attracted public attention. These incidents were not caused by the cryptocurrency itself. Thus, we must exercise caution to ensure there is no adverse effect on the utilization of cryptocurrency, which has great potential.

While writing this executive summary, an alert regarding memcached access controls was issued. A few days later, it was announced that a DDoS attack of 1.35 Tbps--one of the largest ever--had been made against GitHub by exploiting this issue. Previous attack techniques have also maliciously used functions that amplify traffic, such as DNS and NTP. However, memcached has unprecedented destructive power, with an amplification factor of over 10,000. At IIJ, we would like to handle incidents such as these while collaborating with other providers.

IIJ aims to introduce the wide range of technology that we research and develop in this IIR, which is comprised of periodic observation reports that provide an outline of various data we obtain through the daily operation of services, as well as focused research where we examine specific areas of technology.

In Chapter 1, we discuss our SOC Report as the periodic observation report for this volume. As announced in the previous volume, the security report formerly found in this IIR will now be posted in a timelier manner on a website called the wizSafe Security Signal. We will continue to cover security in our periodic observation report once a year. This time we introduce the new Data Analytics Platform that has been renewed to create unique security intelligence at IIJ's SOC, and discuss the activities that were discovered using this platform, based on the details reported in the wizSafe Security Signal over the past six months.

Chapters 2 and 3 are our focused research. First, in Chapter 2 we give an overview of IIJ's full MVNO initiatives. MVNO refers to a mobile virtual network operator, but there are a variety of business models based on the scope of functional elements you will own and operate yourself. A full MVNO uses a business model where you own all functional elements yourself except for the wireless access. Though owning a range of equipment, it is possible to offer services as an MVNO that until now only an MNO (a traditional mobile communications operator that has wireless access) could provide. We explain what it means to be a full MVNO, and what this enables you to do.

In Chapter 3, we discuss the "Hayabusa" open-source software implemented by IIJ Innovation Institute as a system capable of quickly storing and retrieving a large amount of logs output by a range of devices from multiple vendors. In the field of network and system operations, it is necessary to store logs output from hardware and software and display any statistical information or search through the logs for troubleshooting. Log information is also crucial when handling security incidents. Hayabusa was developed as a system to fulfill the need to perform high-speed searches through the vast amount of logs stored by large-scale systems. Here, we introduce the results of tests based on actual ShowNet syslog data collected at Interop Tokyo, along with information about Hayabusa's implementation and future challenges.

IIJ continues to strive towards improving and developing our services daily, while maintaining the stability of the ICT environment. We will keep providing a variety of services and solutions that our customers can take full advantage of as infrastructure for 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.

SOC Report

1.1 Data Analytics Platform Innovation

Since fiscal 2016, IIJ has been working on bolstering its security business from the perspective of human resources, systems, and operations to provide better services to customers. One of the efforts to strengthen our systems is the innovation of our Data Analytics Platform for the comprehensive analysis of security-related information. The Data Analytics Platform aims to cope with increasingly sophisticated cyber attacks by detecting malicious malware activity and other threats through a multifaceted approach to analyzing data. such as IIJ service logs. This enables us to implement preventative measures and post-incident handling for security threats appropriately.

Use of the Data Analytics Platform will make it possible to identify sophisticated cyber attacks that until now have been difficult to detect at an early stage. Here, we will discuss three features of the Data Analytics Platform: (1) Data collection and analysis as an ISP, (2) Data obtained through corporate activities, and (3) Analysis using IIJ's unique big data infrastructure.

■ (1) Data Collection and Analysis as an ISP

In addition to compromising devices by exploiting vulnerabilities, attackers also attempt to launch successful attacks by tricking users into making mistakes. For this reason, it may not be possible to detect cyber attacks using a single security device. Defense in depth using different security devices is an effective way to cope with increasingly complex cyber attacks. Likewise, multi-layer detection equipment provides a robust method for detecting cyber attacks. The Data Analytics Platform performs multifaceted analysis of various logs to detect cyber attacks. The logs collected and analyzed naturally include those from security devices such as firewalls, IPS/IDS, and anti-virus solutions provided as IIJ services, along with those from Web access and incoming/outgoing email. Furthermore, as an ISP we are in the unique position of being able to include logs from backbone traffic and DNS queries in the analysis data. By collecting and analyzing a wide range of logs, we can detect the malicious activity of malware that was previously difficult to uncover.

■ (2) Data Obtained Through Actual Corporate Activities

The Data Analytics Platform detects cyber attacks by utilizing logs from IIJ services, in addition to data from sources such as honeypots and crawlers. IIJ services are used in business activities of companies, so we can utilize data on cyber attacks encountered during these activities. For example, we can gather and analyze logs from the IIJ Secure Web Gateway Service and IIJ Secure MX Service related to user activities for over a million corporate accounts in Japan, including Web access, email, and anti-virus scan result logs. Collecting and analyzing these data enables us to take a close look at the campaign trends of targeted attacks against Japan, or attacks against an unspecified number of users.

■ (3) Analysis Using IIJ's Unique Big Data Infrastructure

Sophisticated cyber attacks can lead to the concealment of malware that continues to exfiltrate confidential data without being discovered by the targeted organization. Even if you are initially unable to avoid being compromised by a 0-Day attack, the speed with which you discover it can determine whether you suffer damages. Figure 1 shows an overview of the Data Analytics Platform architecture. The Data Analytics Platform is constructed as IIJ's own big data infrastructure, based on the open source Hadoop framework. This unique infrastructure makes it possible to separate over hundreds of thousands of logs per second into logical field and value groups, so they can be analyzed in a database using a multifaceted approach. It also enables analysis results to be obtained with sufficient speed when analyzing logs in database form. Furthermore, the architecture is built to allow the expansion

of imported data types and improvements to performance, so it can cope with cyber attacks that continue to become more sophisticated and complex. The Data Analytics Platform makes it possible to identify from a vast amount of data, the malicious activity of threats such as malware at an early stage.

There are various options for output, including reputation information such as a blacklist of C&C servers, and attack observation data. We began posting the attack observation data output on our wizSafe Security Signal*1 security information site from October 2017.

The wizSafe Security Signal publishes periodic observation data each month, while also providing timely security information in blog form.

The following sections discuss the activities that were discovered using the Data Analytics Platform, based on the details reported in the wizSafe Security Signal over the past six months. Of note was a cryptocurrency mining service, DDoS attacks, and attacks targeting vulnerabilities in Apache Struts 2.

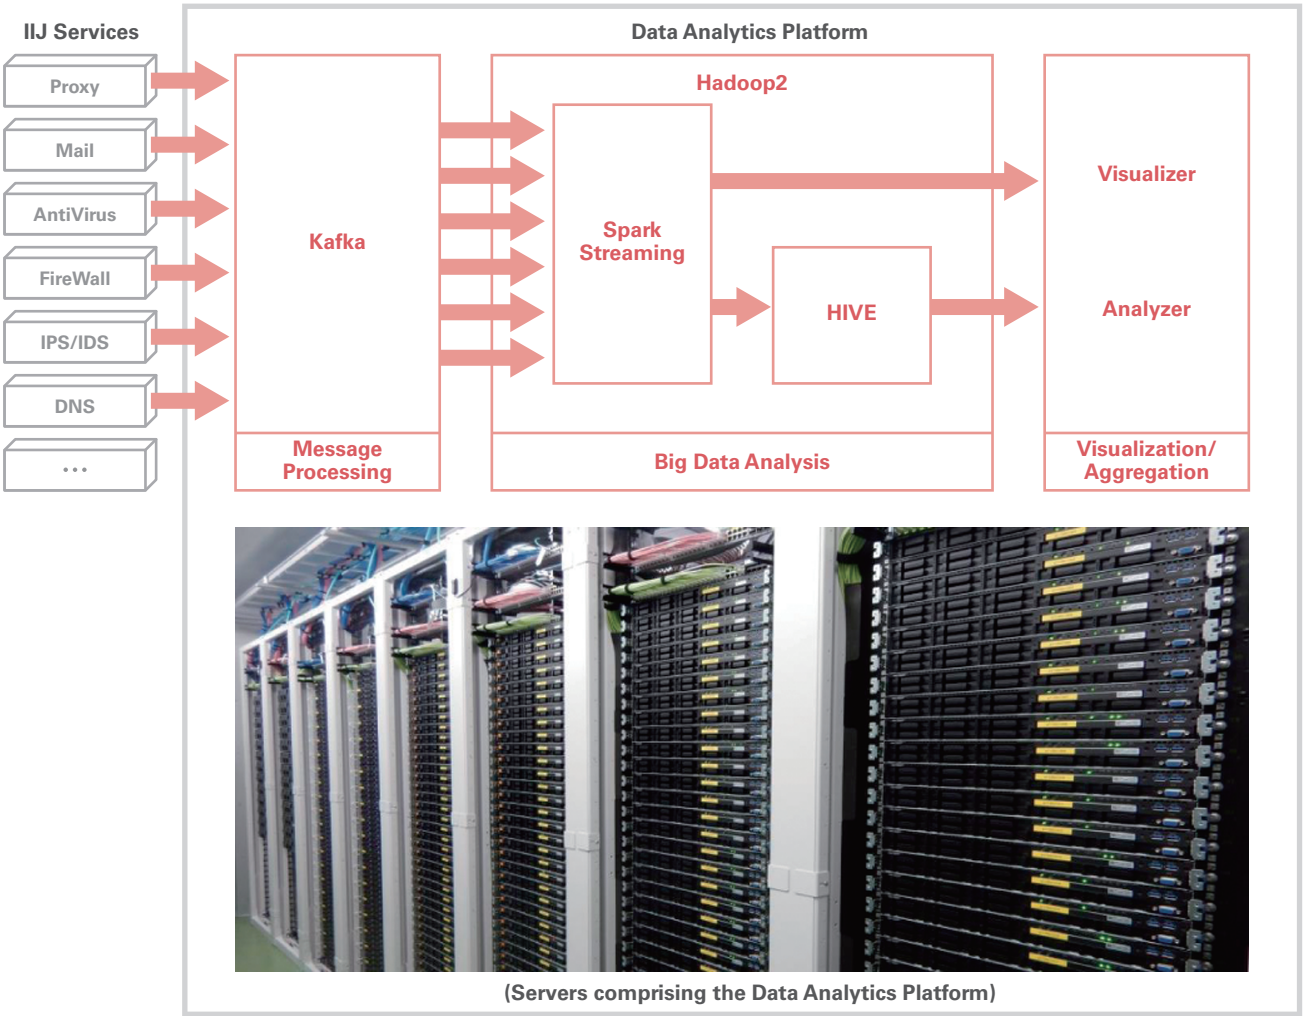

Figure 1: Overview of Data Analytics Platform Architecture

*1 wizSafe Security Signal - Milestones for Security and Safety (https://wizsafe.iij.ad.jp/) (in Japanese).

1.2 Cryptocurrency Mining Services

2017 was also a year where the sharp rise in the price of virtual currencies such as Bitcoin gained considerable attention. For example, the value of Bitcoin increased by about 20 times over the course of the year. Since a few years ago, virtual currencies such as Bitcoin have been used in ransomware demands due to their convenience and the anonymity they provide when sending and receiving them. For example, the Cryptlocker ransomware discovered in 2013 requested Bitcoin as one of the ransom payment methods. Many ransomware attacks demand ransom payments in cryptocurrency, but other than its use in financial transactions, some malware also performs fraudulent mining of said currency. There has also been an incident*2 where malware-infected devices form a botnet that mines for cryptocurrencies.

On September 14, 2017, a cryptocurrency mining service called Coinhive was launched. Coinhive is a service that enables the mining of a cryptocurrency called Monero using JavaScript. A website operator embeds the mining JavaScript in their site to induce mining of the Monero cryptocurrency using the CPU resources of website visitors' PCs. Under the Coinhive system, 70% of mining revenue is distributed to website operators. With cryptocurrency prices skyrocketing, more and more website operators are using the Coinhive service as a revenue source in place of Web advertisements. This kind of cryptocurrency mining service existed before, but since Coinhive's inception, a string of similar services using JavaScript such as Cloudcoins and Coinlab have appeared in quick succession.

Meanwhile, there have also been cases*3 where attackers have adopted cryptocurrency mining services that use JavaScript.

An attacker can obtain revenue from cryptocurrency mined on the PC of a website viewer by modifying the website and embedding JavaScript for mining. Some mining JavaScript consumes a lot of CPU resources, and there are concerns that this could affect the remaining battery life of notebook PCs or smartphones. This demonstrates that cases involving the fraudulent use of cryptocurrency mining services are on the rise, and some anti-virus software now treat JavaScript for mining as a risk and blocks it.

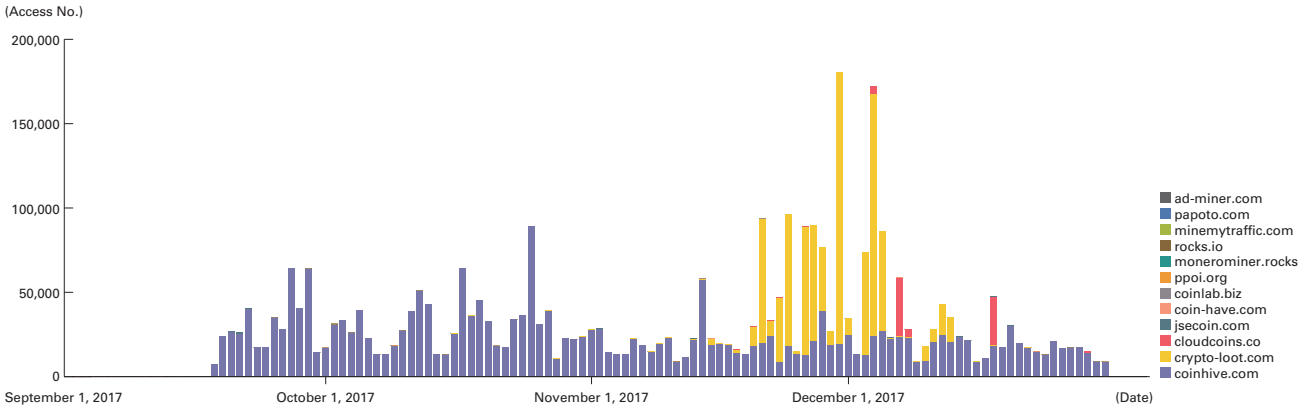

Figure 2: Access to Cryptocurrency Mining Services in 2nd Half of 2017

*2 Kaspersky Lab Daily, "Got any hidden miners? I wouldn't be so sure..." (https://www.kaspersky.com/blog/hidden-miners-botnet-threat/18488/).
*3 wizSafe Security Signal, "Cases of Cryptocurrency Mining Scripts Embedded Through Website Alteration" (https://wizsafe.iij.ad.jp/2017/10/94/) (in Japanese).

Figure 2 shows access numbers for the cryptocurrency mining service observed over the IIJ Secure Web Gateway Service since September 2017, when the Coinhive service launched. We can see that there are an increasing number of different cryptocurrency mining services, and access numbers for these services climbed until early December. Because this data is from the website access of enterprise users, access numbers are relatively low on Saturdays and Sundays, public holidays, and at the end of the year when corporate users were not active.

Whether the use of cryptocurrency mining services remains appealing in the future will depend on how the prices for cryptocurrencies change. The ease of using these services will also be affected by the risk assessment of JavaScript for cryptocurrency mining services by website viewers. For example, the operator of a legitimate website would not want to embed JavaScript that is blocked by anti-virus software on the PC of website visitors.

Cryptocurrency prices and security vendor trends also affect the ability of attackers to use cryptocurrency mining services. If the use of cryptocurrency mining services loses its appeal, attackers will likely look for other ways to obtain money.

1.3 DDoS Attacks

During the second half of 2017 (July to December), IIJ observed an average of 20.8 DDoS attacks per day and 638 DDoS attacks per month, totaling 3,828 attacks overall. Figure 3 shows the number of DDoS attacks that occurred in the second half of 2017. Because September contains dates such as September 18 when the Manchurian Incident started, we took precautions against the occurrence of many DDoS attacks, but we observed no related attacks to this in 2017. DDoS attacks have become a daily occurrence, and we believe fluctuations in the number that take place on a given day are affected less and less by specific time periods.

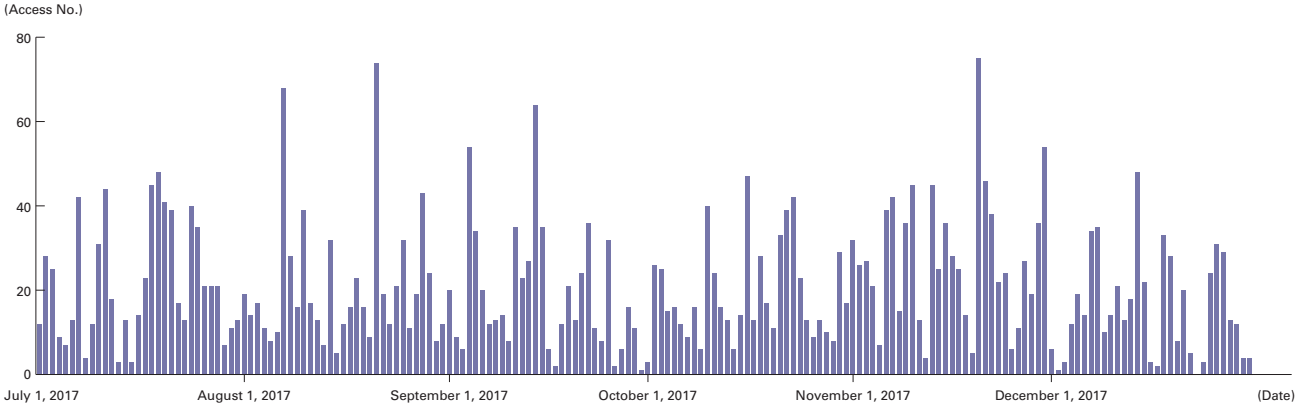

Figure 3: Trends in DDoS Attacks in 2nd Half of 2017

The largest DDoS attack observed during the survey period was observed in October 2017 and generated 16.55 Gbps of traffic using up to 1.79 million pps packets. The main attack technique used was a UDP flood, which causes line congestion. The longest sustained attack was a DDoS attack observed in September 2017 that continued for 46 hours 57 minutes. This attack mainly used the ICMP protocol. Table 1 shows the number of DDoS attacks that occurred each month in the second half of 2017, along with the largest and longest attacks and the methods used for each.

Threatening emails demanding the payment of money continued to be distributed during this observation period. In some cases, a small DDoS attack was carried out before threatening to launch a larger one. Since around December 2017, there were multiple incidents where notice of DDoS attacks was provided on Twitter with targets specified in Japanese. These were followed by multiple DDoS attacks being launched against the targeted websites, causing a denial of service. Targets included the websites of government and municipal offices.

There are various types of DDoS attacks, such as those aimed at monetary theft, and those carried out as hacktivist activity. However, in each case, attacks are usually intended to cause a big impact on business. Figure 4 shows the occurrence rates for DDoS attacks by day and time period. There were fewer attacks at night and on weekends than in the daytime and on weekdays. Business hours from 7:00 to 13:00 had more than triple the occurrence rate of the early morning hours of 1:00 to 7:00. There was also a difference of more than three times the number of incidents, when comparing Saturday and Monday.

Table 1: Max DDoS Attack Scale / Length in 2nd Half of 2017

Month / Year	Incidents	Max Attack Scale / Method		Max Attack Length / Method	
July 2017	673	17.86 Gbps	NTP Amplification	5 hrs 18 mins	IP Fragmentation/UDP
August 2017	655	10.16 Gbps	IP Fragmentation/UDP	16 hrs 43 mins	Flood
September 2017	575	12.06 Gbps	NTP Amplification	46 hrs 57 mins	ICMP
October 2017	593	16.55 Gbps	UDP Flood	24 hrs 9 mins	IP Fragmentation/UDP
November 2017	843	8.93 Gbps	IP Fragmentation/UDP	6 hrs 23 mins	UDP Flood
December 2017	489	13.39 Gbps	DNS Amplification	14 hrs 22 mins	ICMP/DNS Amplification

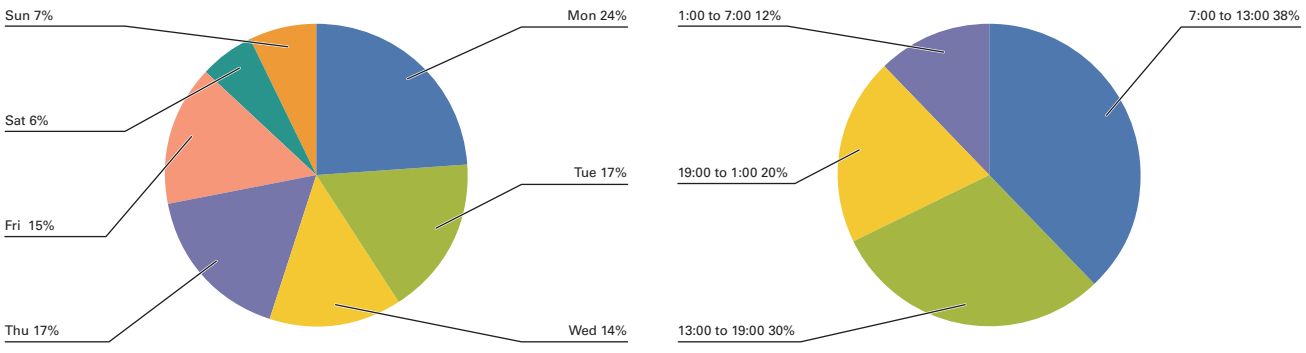

Figure 4: DDoS Attacks by Day / Time of Day

1.4 Attacks Targeting Struts 2 Vulnerabilities

In 2017, attacks targeting critical vulnerabilities in Apache Struts 2 occurred frequently. As shown in Table 2, six vulnerabilities disclosed for Apache Struts 2 in 2017 possibly allowed remote command execution.

Of these vulnerabilities, CVE-2017-5638 (S2-045/S2-046) disclosed in March 2017 could be exploited in attacks relatively easy. Because the environmental requirements for attack targets were comparatively loose, unless a security patch was applied, the success rate for an attack was extremely high. As a result, there was a rash of cyber attacks that led to damages such as information leaks immediately after the vulnerability was disclosed. Table 3 provides examples of security incidents where Apache Struts 2

Table 2: Remote Command Execution Vulnerabilities in Apache Struts 2 (2017)

Bulletin#	Description	CVE#	CVSS v3 Base Score
S2-045	Possible Remote Code Execution when performing file upload based on Jakarta Multipart parser.	CVE-2017-5638	10 Critical
S2-046	Possible RCE when performing file upload based on Jakarta Multipart parser (similar to S2-045)	CVE-2017-5638	10 Critical
S2-048	Possible RCE in the Struts Showcase app in the Struts 1 plugin example in Struts 2.3.x series	CVE-2017-9791	9.8 Critical
S2-052	Possible Remote Code Execution attack when using the Struts REST plugin with XStream handler to handle XML payloads	CVE-2017-9805	8.1 High
S2-053	A possible Remote Code Execution attack when using an unintentional expression in Freemarker tag instead of string literals	CVE-2017-12611	9.8 Critical
S2-055	A RCE vulnerability in the Jackson JSON library	CVE-2017-7525	8.1 High

Table 3: Examples of Security Incidents Targeting Struts 2

Month/Year Disclosed	Incident Overview
March 2017	Leak of 676,290 pieces of information from a credit card payment site for metropolitan tax*4
April 2017	Leak of up to 23,000 pieces of information from a map data site*5
May 2017	Unauthorized access to a public server at the National Institute of Information and Communications Technology (NICT)*6
June 2017	Leak of up to 4,335 pieces of information from the Land General Information System*7
September 2017	Leak of information on up to 143 million individuals from Equifax*8

*4 Tokyo Metropolitan Government, "Unauthorized access to credit card payment site for metropolitan tax" (http://www.metro.tokyo.jp/tosei/hodohappyo/press/2017/03/13/02.html) (in Japanese).
*5 Ministry of Internal Affairs and Communications, "Unauthorized access to map-based small-area analysis (jSTAT MAP)" (http://www.soumu.go.jp/menu_news/s-news/01toukei09_01000023.html) (in Japanese).
*6 National Institute of Information and Communications Technology, "Unauthorized access exploiting Apache Struts 2 vulnerability" (https://www.nict.go.jp/info/topics/2017/05/170502-1.html) (in Japanese).
*7 Ministry of Land, Infrastructure, Transport and Tourism, "About the possibility of unauthorized access and information leaks in the Land General Information System" (http://www.mlit.go.jp/report/press/totikensangyo05_hh_000129.html) (in Japanese).
*8 Equifax, "Equifax Releases Details on Cybersecurity Incident, Announces Personnel Changes" (https://investor.equifax.com/news-and-events/news/2017/09-15-2017-224018832).

vulnerabilities were exploited. It has been revealed that the CVE-2017-5638 (S2-045/S2-046) vulnerability was exploited in some incidents. Equifax, a major credit information company in the United States, also announced a leak of personal information caused by CVE-2017-5638 (S2-045/S2-046) in September 2017. It is thought that this incident could have resulted in the leak of the personal information of 143 million customers in the United States, Canada, and the UK.

Figure 5 shows the number of attacks per site that exploited Apache Struts 2 vulnerabilities observed by IIJ in the second half of 2017 (July to December). We can see that attacks against CVE-2017-5638 (S2-045/S2-046) disclosed in March 2017 account for the majority. Some attacks targeting CVE-2017-5638 (S2-045/S2-046) are estimated to be large-scale explorative attacks made with the same tool using a botnet. The rapid increase in incidents*9 from October 20, 2017, is due to these attacks.

It has already been more than six months since CVE-2017-5638 (S2-045/S2-046) was disclosed, but there are still many attack targets on the Internet that can be exploited, and attackers still have their sights set on these Web servers.

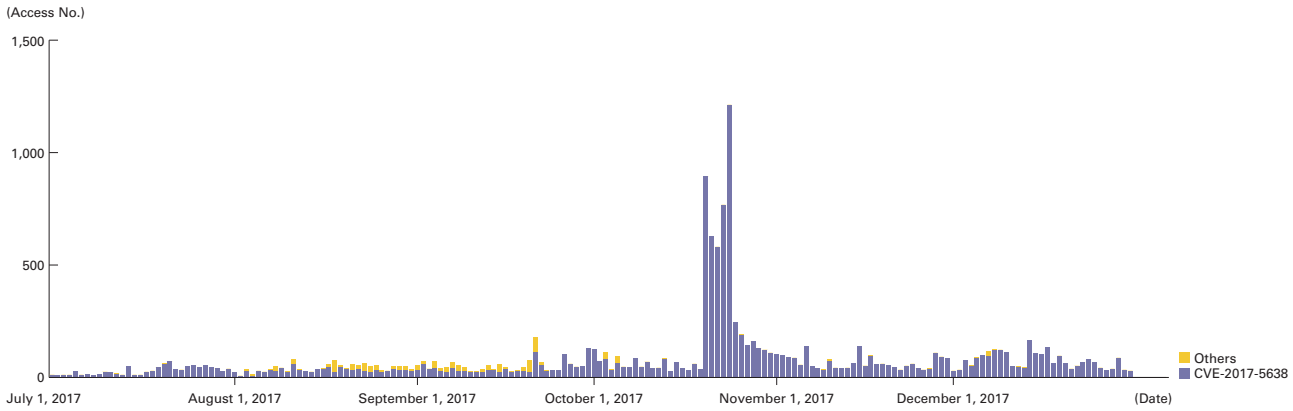

Figure 5: Attacks Per Site Targeting Apache Struts 2 in 2nd Half of 2017

*9 wizSafe Security Signal, "Observations of Attacks Targeting Apache Struts 2 Vulnerabilities" (https://wizsafe.iij.ad.jp/2017/11/106/) (in Japanese).

Apache Struts 2 is used as a Web application framework in various websites. Websites that did not apply Apache Struts 2 security patches properly were affected by attacks. Apache Struts 2 was the target of many attacks in 2017, but all servers and software that can be accessed via the Internet are considered potential targets by attackers. If even a single server or software vulnerability is not properly managed, it could lead to damages such as information leaks. It is essential to regularly review whether proper vulnerability management is being performed, regardless of whether Apache Struts 2 is used.

1.5 Conclusion

In this report, we looked at some of the results of long-term ongoing analysis of trends in cyber attacks utilizing the Data Analytics Platform. We have only scratched the surface of what the new Data Analytics Platform is capable of. Going forward, IIJ will continue to increase the type and amount of data incorporated and introduce new analytical techniques such as machine learning and AI, so we can take immediate action against sophisticated cyber attacks.

Author:
Kiyoshi Saito
General Manager, Security Business Department, Advanced Security Division, IIJ
As Director of the Security Business Department, Mr. Saito is engaged in the operation, implementation, and support of managed security services, as well as security consulting, SI, SOC integration, and business development.

Author:
Tsutomu Nakajima
Deputy Manager, Security Operations Center, Security Business Department, Advanced Security Division, IIJ
Mr. Nakajima leads a team of security analysts as Deputy Director of the SOC. He is also engaged in big data analysis work as an analyst himself, including incident analysis and the creation of security intelligence.

Why IIJ Seeks to Become a Full MVNO

2.1 MVNO Business Models

In 2018, IIJ finally begins full MVNO services, the biggest challenge we have had since launching our MVNO business in 2008. The phrase “full MVNO” is still not used all that commonly in Japan, and it may be difficult to understand it correctly. However, tens of MVNOs around the world have already successfully transformed their business model to “full MVNO,” and are providing advanced and diverse services through their platform.

“Full MVNO” is a phrase that defines a particular MVNO business model. Figure 1*1*2 shows typical MVNO business models, and the key to their classification lies in how many of the elements the MVNO operates by themselves (Figure 2*3).

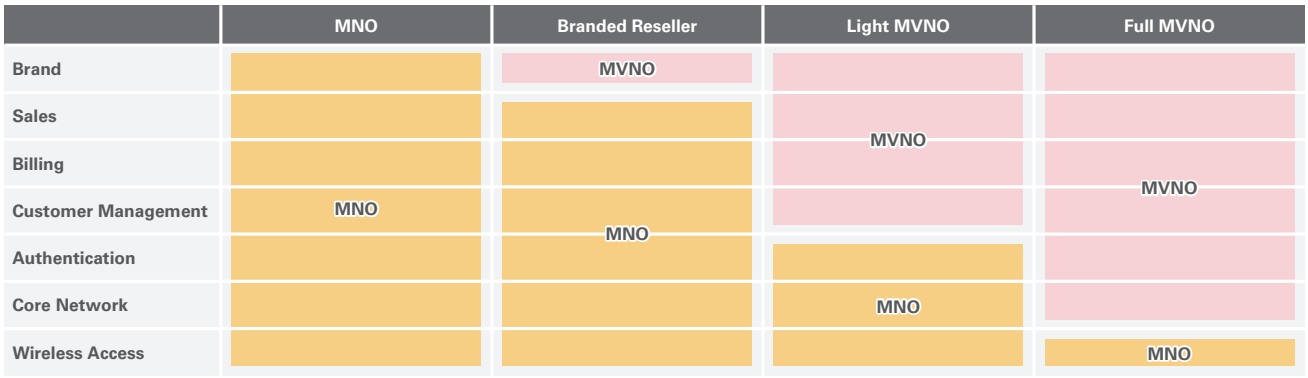

Figure 1: Types of MVNO

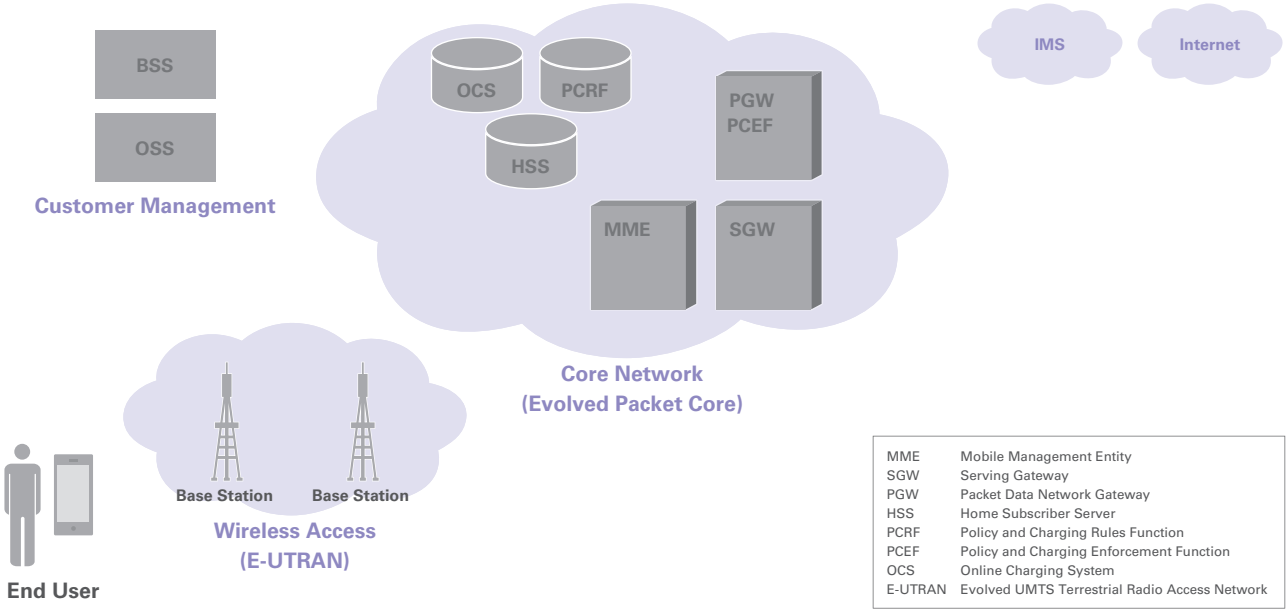

Figure 2: 4G LTE Mobile Communications Network Overview

*1 Figure 1 shows various types, but actual MVNOs may not always fit one of these categories. Many MVNOs could settle between these categories due to market trends and regulations in each country, as well as relationships between MNOs and MVNOs. There are also local varieties such as the Layer 3 MVNOs and Layer 2 MVNOs in Japan.

*2 Figure 1 is not intended to indicate which one is superior or inferior. Be aware that it is important to select a favorable business model based on a company’s characteristics and business objectives.

*3 This shows typical 4G LTE nodes, but MVNO business models differ as to which parts other than the wireless access are operated by the MVNO themselves and which parts are provided by their host MNO.

As the name indicates, an MVNO (Mobile Virtual Network Operator) is a virtual provider that essentially offers services under its own brand by relying on equipment that they do not operate themselves. However, when you are reliant on other companies for various elements other than your brand, you must be concerned about your business lacking uniqueness, making differentiation difficult. For that reason, depending on the business objectives of each company, there is the option of directly operating some, if not all, elements and equipment yourself, without relying on other companies.

For example, for a company like Disney that has a very strong brand and is highly competitive even when relying on the host MNO for all other elements, a "Branded Reseller" business model is probably the most suitable. The "Light MVNO" model that performs sales and customer management themselves while relying on the host MNO for network and authentication equipment is also spreading as a standard business model both in Japan and across the globe. "Full MVNO" is the MVNO business model closest to an MNO, operating most of the equipment other than wireless access by themselves.

2.2 Full MVNOs: Their Competitive Edge and Hurdles

The competitive edge of full MVNOs lies in the network equipment they possess, or their core network*4, as well as their authentication facilities. This equipment can be served by an MNO when using another MVNO business model, but by operating them by yourself it is easier to maintain the uniqueness of your business and differentiate yourself from other MVNOs or MNOs. Not only that, but by operating the core network at the heart of mobile communications, you can transition from a light MVNO business model where you are bound to a designated MVNO agreement with a single MNO, and take on new business areas through collaboration with multiple MNOs and MVNOs.

On the other hand, a full MVNO has the highest cost hurdle due to the investment in equipment and the many human resources required to operate it. Consequently, there seem to be many cases around the world where low value-added services like the low price plans for smartphones are provided by light MVNOs, and providers entering new business areas with high added value like IoT, international services, security, and FinTech go the full MVNO route.

2.3 The Shift Towards Full MVNOs

Until now, there were no full MVNOs in Japan, so IIJ is the first to commercialize it. Internationally, there are still only a limited number of countries where full MVNOs exist, including some European nations. For many countries the commercialization of full MVNOs will be a future challenge. As mentioned above, a full MVNO requires significant costs, so we have to consider whether an MVNO with the scale to make this investment already exists. MNOs must also have the maturity to accommodate partnerships with full MVNOs. In some cases, national telecommunications policies or regulations may not allow full MVNOs, and it may be necessary to hold discussions.

That said, looking at future advancements in mobile communications such as IoT and 5G, the innovation and diversity that full MVNOs bring to the mobile communication market will be crucial for all countries. We believe that for many nations the introduction of full MVNOs to the market will become a practical policy issue sooner or later. In Japan, discussions held at the Information and Communications Council's 2020-ICT Basic Policy Special Committee in 2014 recognized that full MVNOs were an important policy issue, and talks between the providers, IIJ, and NTT Docomo gained momentum. We expect that trends like this will continue to spread around the world.

*4 Core Network: A network that transports signals (signaling) for connecting devices to a network and data (user data) related to calls and communications within a mobile communications network. In 2G/3G architecture, legacy protocols are still in use. But nowadays, for the purpose of maintainability those legacy protocols are often transported by IP. In 4G architecture, all protocols are based on IP.

2.4 Full MVNOs and the Unbundling of HLR/HSS

Discussions related to full MVNOs in Japan, including those at the 2020-ICT Basic Policy Special Committee, have centered on the unbundling of HLR*5 and HSS*6 for MVNOs. HLR and HSS are nodes that have important roles in the core network for mobile communications. They are often described as subscriber management devices, offering the following diverse range of functions:

1. A database for storing and managing information such as the MSISDN (phone number) or IMSI (international number for subscriber identification) of all subscribers.
2. Responds to queries from each node within a network such as the MME, and authenticates the connection and function availability of subscribers.
3. Stores information on networks and switching equipment serving devices in real time to enable services such as voice calls and SMS (location management).
4. Stores information corresponding with SIM cards, and provides the encryption functions (along with other nodes) necessary for keeping the content of communications secret.

Until now, HLR and HSS had been operated only by MNOs in Japan, and were not available from MVNOs. IIJ is the first MVNO in Japan that operates HLR and HSS independently and connects them to an MNO's network directly*7.

However, care must be taken to treat the operation of HLR and HSS as the definition of a full MVNO, despite there being various nodes in a core network. The wider network model of full MVNOs that operate other nodes in core networks such as MME*8 and SGW*9 is also discussed worldwide. This is also called "RAN*10 sharing" between MNOs and MVNOs, but IIJ's commercial implementation of full MVNO does not adopt this wider definition. MME and SGW are almost incapable of developing high value-added services by themselves. However, for the upcoming 5G technology, introducing added value such as network slicing or NFV to core networks is a major topic, and we consider this to be a subject for future analysis.

2.5 MNC and SIMs

The benefits of an MVNO operating HLR and HSS are not always the direct result of HLR/HSS functions. One example is MNC*11. MNCs are the numbers that serve as identifiers for mobile communications networks when they are connected to one another. This is similar to an AS number in an IP network. Like IP networks, connections are made between mobile communications networks on a global scale, so the numbers used as identifiers must be globally unique. In Japan, an MNC is a five-digit number allocated by the Ministry of Internal Affairs and Communications under the Regulations for Telecommunications Numbers (Ministerial Ordinance), based on ITU-T*12 recommendation E.212. The three digits at the head of it are a country code called the MCC, for which Japan has been assigned 440 and 441.

*5 HLR: An abbreviation of Home Location Register. A node that stores information on subscribers and manages location information on 2G/3G mobile networks.
*6 HSS: An abbreviation of Home Subscriber Server. A 4G LTE network and IMS node that provides functions similar to HLR.
*7 Excluding "MVNOs that are also MNOs," such as KDDI and Softbank in Japan.
*8 MME: An abbreviation of Mobile Management Entity. A node that handles signaling in 4G core networks.
*9 SGW: An abbreviation of Serving Gateway. A node that handles user data in 4G core networks.
*10 RAN: An abbreviation of Radio Access Network. A network for connecting base stations (wireless station equipment) to core networks in mobile communications. For 3G this is referred to as UTRAN (UMTS Terrestrial Radio Access Network), and for 4G it is referred to as E-UTRAN (Evolved UTRAN).
*11 MNC: An abbreviation of Mobile Network Code. A unique number used to identify mobile communications networks.
*12 ITU-T: The Telecommunication Standardization Sector of the International Telecommunication Union.

According to the Regulations for Telecommunications Numbers, the requirement for allocation of this MNC is "the installation of equipment for identifying devices." IIJ has met this requirement by installing and operating HLR and HSS independently, and we have been allocated the MNC of "44003" by the Ministry of Internal Affairs and Communications, which is a first for an MVNO in Japan. Using this MNC, IIJ can connect its own HLR and HSS to DOCOMO and other mobile communications networks, making it possible to implement roaming services.

Another benefit is the capability to issue our own SIMs. SIMs are IC cards that are issued to each subscriber. They hold a number for subscriber identification (IMSI*13) and an encryption key. Subscribers can access the services of a mobile communications network by inserting this SIM into a device. MVNOs in Japan also issue SIMs and offer them to subscribers, but this is merely passing on SIMs provided by their host MNO. Now IIJ has obtained its MNC, we can allocate IMSI to the customer using our own equipment, and we can issue our own SIMs independently*14.

In terms of technology, there are two advantages when an MVNO is able to provision SIMs. One is greater flexibility to provide services. In the past, there was no way to provision SIMs without relying on the system of the host MNO. That means it was not possible to offer services the MNO's system did not allow. For example, under conventional light MVNO schemes, once a SIM is provisioned and its use started, it continues to be available until it is deactivated, and after deactivation, the SIM cannot be reused. That means it is difficult to activate it temporarily during pre-shipment inspections in cases where a communications module is to be embedded into the product for IoT purposes. In other words, because the process of removing the SIM used in the pre-shipment inspection and inserting a new SIM is performed after the check, the inspection could lose its meaning. Limitations like this are coming from BSS or OSS*15 on the MNO side. Ideally, if an MNO could develop BSS or OSS with enough flexibility to meet the demands of MVNOs, even light MVNOs would be able to overcome this issue. However, in reality the barriers are quite high.

IIJ's full MVNO services support the temporary activation of SIMs and the recycling of deactivated SIMs through IIJ's BSS and OSS implementations. We believe the merits can be leveraged where flexible activation or deactivation of the SIM is required, such as with IoT.

The other advantage is support for new SIM technology, such as eSIMs*16. In advanced IoT and roaming use cases, the restrictions of plastic SIMs that have been around for thirty years have become a new problem. We expect to see the spread of eSIMs that enable SIM profiles to be downloaded online without removing and inserting them, and virtual SIMs (software SIMs*17) that handle the essentials of SIMs virtually without any physical media or devices. Commercial devices equipped with eSIMs and devices fitted with software SIMs that provide prepaid inexpensive roaming have already appeared. By having a platform that enables independent provisioning of SIMs, we will be able to obtain the benefit of innovations using these new technologies.

*13 IMSI: An abbreviation of International Mobile Subscriber Identity. This is an identifier allocated to each mobile communications network subscriber that is stored on the SIM. The first few digits of an IMSI are the MNC of the issuing operator. HLR and HSS are used to manage the IMSI and handle SIM authentication.

*14 Provisioning: To make preparations for providing network resources to telecommunications service subscribers. In this case, it refers to preparing to provide a mobile communications service by writing the necessary information to the SIM and coupled information to HLR or HSS.

*15 BSS/OSS: An abbreviation of Business Support System / Operation Support System. These refer to systems in the back offices of telecommunications carriers. BSS covers customer-related systems such as those for subscription and billing management, while OSS covers operation-related systems such as those for SIM provisioning and logistics management.

*16 eSIM: A SIM that enables the telecommunications carrier profile stored inside to be rewritten remotely. Physically, a conventional plastic SIM or a dedicated solderable chip will be used. Currently, the GSMA, an industry group for mobile telecommunications carriers, is moving ahead with standardization on two tracks--one for Machine-to-Machine communication and the other consumer devices.

*17 Software SIM: This refers to remotely rewritable SIMs that do not use dedicated hardware, and are instead implemented through software using a processor's trusted execution environment (TEE). It is not a standardized technology, but certain software SIM platforms are compatible with standardized eSIM provisioning systems.

2.6 The Benefits of IIJ's Full MVNO

Some MVNOs in Japan other than IIJ have begun to implement HLR and HSS and secured the flexibility of SIM provisioning. This is some particular preamble of the trend of IoT. So, what are the differences between IIJ's full MVNO scheme and other businesses developing with a similar business model?

We think that the HLR/HSS implemented by other providers will generally only be connected to networks of overseas providers and not MNOs in Japan. Despite this, they can provide the connectivities of other mobile communications networks and can obtain the capability of SIM provisioning as described in this report in the same way as IIJ's full MVNO scheme does. However, in the case of these providers, the connectivity of MNO networks in Japan is provided based on a roaming agreement between the overseas operator and domestic MNO just as with countries other than Japan. As a result, they will only be able to offer services with less flexibility of pricing due to the comparatively expensive roaming charges in Japan.

Even for IIJ, there is no significant difference in countries other than Japan, because roaming services are provided at a relatively high cost via the networks of overseas roaming providers and their roaming agreements. However, in Japan IIJ connects its HLR and HSS to NTT DOCOMO directly, so the domestic data is more reasonable compared to roaming charges. That means we can provide domestic services with highly flexible pricing at a lower cost than our competitors in Japan, which is a major distinguishing characteristic.

Next, we will examine roaming services in countries other than Japan in detail. Generally, payments for international roaming are offset based on mutual use when services are provided bilaterally between two operators with a roaming agreement (each company opens their network to roaming from the other party). Consequently, depending on the roaming fees (tariff) set and the amount of traffic on both sides, it may not be too difficult for a company to procure data connectivities from other countries at levels that are similar to their own network costs. However, when connecting only HLR and HSS to the networks of overseas roaming providers, it is not possible to sell data connectivities in Japan to overseas operators, so they cannot make a bilateral roaming ecosystem. Thus, the cost of procuring overseas data connectivity depends solely on the roaming fees (tariff) set by overseas operators, making it difficult to obtain the flexibility of pricing.

Contrarily, IIJ can sell Japanese data connectivity procured from NTT DOCOMO at the MVNO data rate, to overseas operators through our own HLR and HSS. This means that, although it is necessary to use Multi IMSI*18 technology, we have an environment that enables us to provide bilateral services similar to roaming business between MNOs. We are currently working with overseas telecommunications carriers to discuss the possibility of offering such bilateral services and providing data connectivity in Japan using IIJ's IMSI. We are making preparations to be able to provide overseas data communication services with highly flexible pricing to meet the needs of Japanese customers. This is a business model that is only possible for IIJ, as the only Japanese MVNO with HLR and HSS connections to the host MNO in Japan, and I think that many customers will appreciate the pricing benefits this will provide.

*18 Multi IMSI: Usually a single IMSI is recorded in a SIM, but this technology enables users to switch between different telecommunications providers with a single SIM by changing the IMSI. Unlike the multi-profiles enabled by eSIM, an applet in the SIM handles switching the IMSI.

2.7 Future Challenges

IIJ has paved the way forward by becoming Japan's first full MVNO. This type of business model is one-of-a-kind in Japan, and has only been commercialized in a limited number of countries and regions around the world. For that reason, IIJ will avoid taking further risks, progressing with deployment in phases while verifying the operation of new facilities and systems to be implemented.

Also, services such as voice calls will be out of scope for this full MVNO. With regard to voice calls, we have taken into consideration the strict regulations, such as maintaining call quality, MNP (mobile number portability), and emergency calls. Currently, IIJ provides low-priced SIMs for smartphones as one of our main lines of business, by offering them through our own IIJmio brand and the brands of other partner companies. However, regarding voice calls, we will continue with our light MVNO-based scheme in the future.

IIJ will consider a full MVNO that includes voice call services while fully taking into account the trend of requirements for such a service and its future prospects.

Furthermore, as a full MVNO, we are proactively thinking about new trends such as IoT and 5G. In order to introduce cellular LPWA*19 on a worldwide scale, and enable a variety of added value through network slicing and 5G, we are considering what kind of equipment we should construct with our host MNO, and which advanced and diverse services we should forge ahead with as an MVNO. IIJ believes that as a front-runner in the MVNO space, we have a great responsibility to the MVNO industry going forward. We remain committed to the belief that an MVNO is not merely a second-tier provider piggybacking on the MNO, but a provider based on a unique business model that enables a variety of services not provided by MNOs.

Author:
Futoshi Sasaki
Deputy General Manager, Strategy and Business Development, MVNO, IIJ.
Since joining IIJ in 2000, Mr. Sasaki has been engaged in the operation, development, and planning of network services.
He was one of the founding members of IIJ's MVNO project in 2007 and has been in charge of corporate and personal MVNO services ever since.
He is a member of the MVNO Committee of the Telecom Services Association, an MVNO industry group.

*19 Cellular LPWA: A Low Power Wide Area mobile communications service for IoT that uses mobile phone technology, and is provided at a frequency (licensed band) that requires licenses to be obtained. Cellular LPWA communications standards such as LTE-M and NB-IoT are expected to be commercialized in 2018 and beyond.

Hayabusa: Simple and Fast Full-Text Search Engine for Massive System Log Data

3.1 Background and Objectives

In the field of network and system operations, network administrators store logs to display statistical information and/or search the logs to identify the source of trouble, to ensure stable network operations and perform troubleshooting. When responding to security incidents, as with troubleshooting, there are also cases where logs are used to investigate the type of incident that occurred.

In large-scale networks, communication logs are output in bulk from many servers, networking and security devices on a daily basis. Network administrators use systems to accumulate these logs in storage systems so that they can search the logs quickly. Large-scale search and storage systems tend to require the use of clustering systems or dedicated management software, and in many cases, operators have to use their time to manage these systems which should be spent for analyzing logs.

If we can build a system to “store” and “search” logs simply without the use of complex clustering systems, network administrators would no longer need to devote as much time to managing the storage and search systems, allowing them to focus more on their core tasks of network troubleshooting and the analysis of security incidents.

In this report, we discuss the Hayabusa open-source software that implements a system capable of quickly storing and retrieving a large amount of logs output by a wide range of devices from multiple vendors. We also introduce a conceptual model for a distributed system using the same software that dramatically improves search speed by scaling out the system’s searching capability when the volume of logs handled by the system increases.

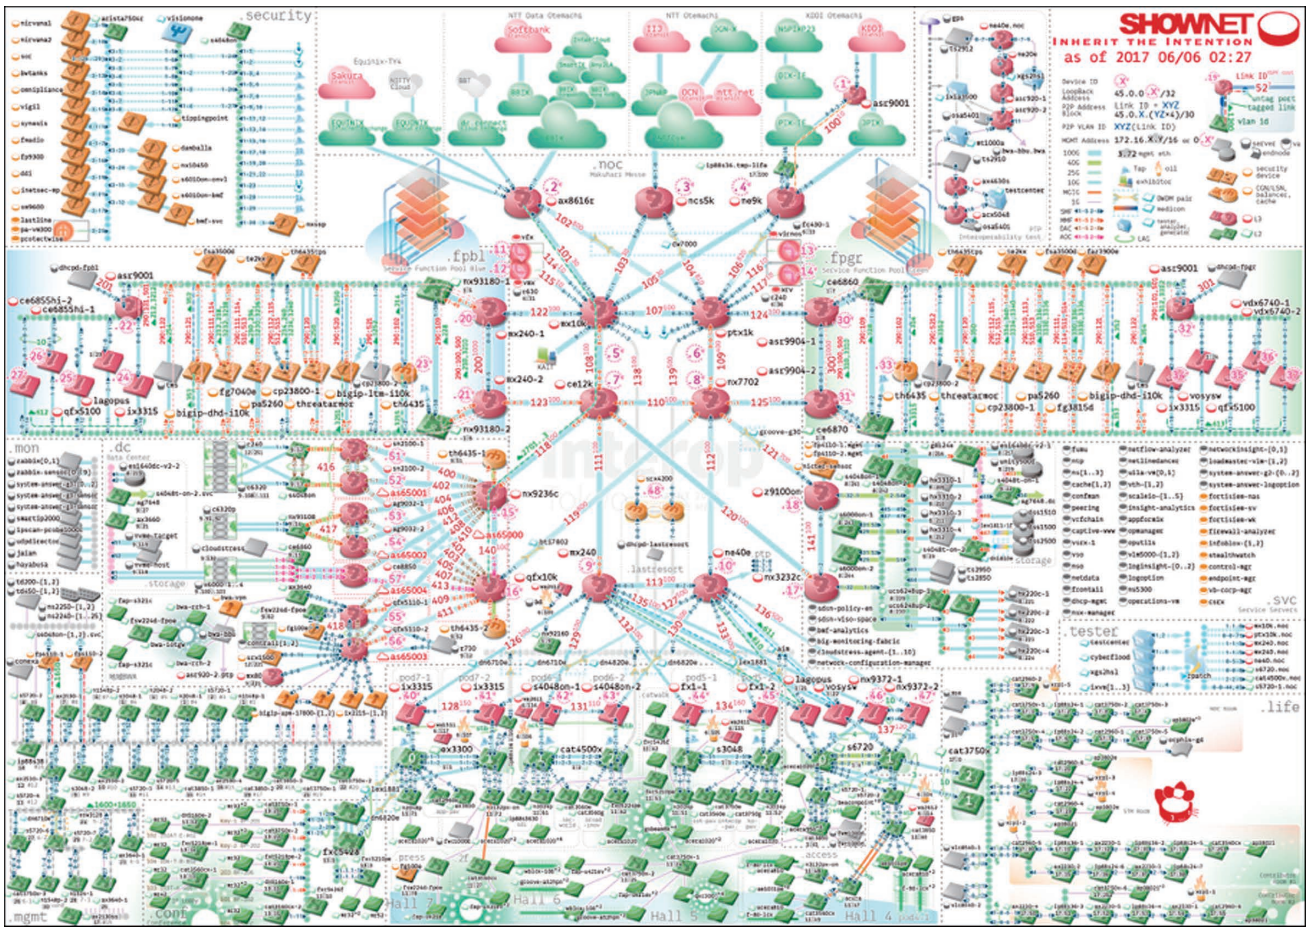

Figure 1: 2017 ShowNet Topology Diagram

To evaluate the search speed of Hayabusa using a large amount of log data output by various devices from multiple vendors, we performed tests based on actual ShowNet syslog data collected at Interop Tokyo*1, output from over 600 servers, networking and security devices.

3.2 ShowNet

ShowNet is a test network environment for verifying interoperability and performing demonstrations constructed at Interop Tokyo, which is held annually at Makuhari Messe. Because it also offers an Internet access service to exhibitors, it has two-fold aspects, one for a real business network service and the other for an experimental network service.

As shown in Figure 1, ShowNet is comprised of hundreds of pieces of equipment, including not only the products in the market but also many products that are connecting to a network on which services are running for the first time, with some devices and software still in the prototype stage.

NOC (Network Operation Center) members operating ShowNet combine these devices and software to design, build, and operate networks to provide services. To construct a stable system, NOC members assess system bugs and errors along with the feasibility of configurations by collecting and analyzing syslog data output from the various equipment. Since ShowNet is comprised from various equipment incorporating a number of cutting-edge software and hardware equipment, so it is difficult for operational members managing the equipment and software to know what kind of logs that the various equipment will output in advance.

In ShowNet, we run a system to collect and analyze syslog data as one of the monitoring systems. Almost all physical and virtual devices send syslog data to the constructed system with which we can aggregate and search the data. In the past, there have been times when over 20,000 log entries were exchanged in a second, and over 200 million log entries were stored in a single day.

It is possible to infer the format of logs if they come from a small set of specific software or hardware devices, but there are very few cases where the same software or hardware is used in the construction of ShowNet. A large number of messages are also output as debug information for the newest firmware or software. This makes it very difficult to display statistical data and search in real time using standard log storage and analysis software.

3.3 Hayabusa

Hayabusa*2 was designed as a system for high-speed searching through a large amount of syslog data collected at ShowNet. Figure 2 shows the architecture of Hayabusa.

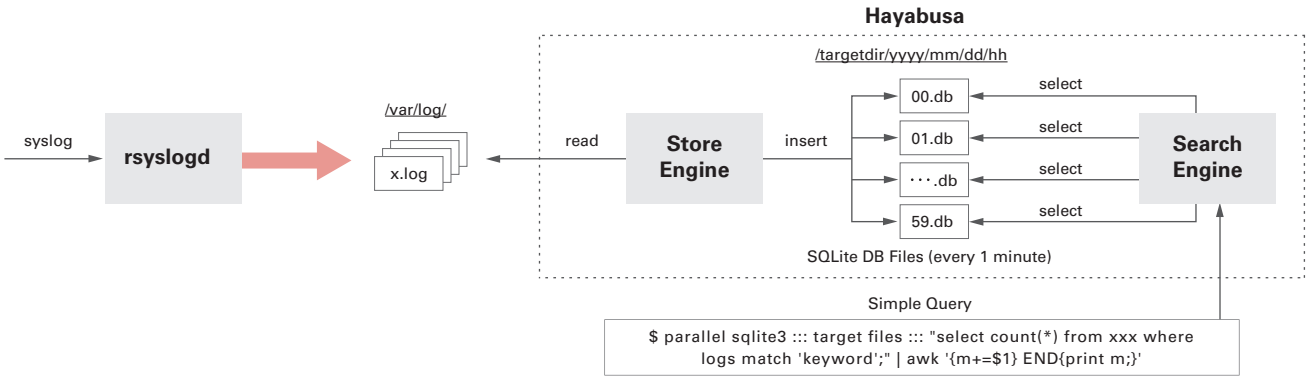

Figure 2: Hayabusa Architecture

*1 Interop Tokyo (https://www.interop.jp/).
*2 Hayabusa (https://github.com/hirolovesbeer/hayabusa).

Hayabusa runs on a stand-alone server and uses multiple CPU cores efficiently to achieve high-speed parallel search processing. It is comprised of two core parts: the StoreEngine, and the SearchEngine. The StoreEngine is launched every minute as a cron job. It opens the target log file and converts the log messages into SQLite3[*3] files. Log data is split into SQLite3 files every minute, and multiple search processes use these files in parallel. The directory for storing logs is defined using a hierarchy that indicates the time as shown below.

/targetdir/yyyy/mm/dd/hh/SQLite3 files for each minute

By defining in this way, it is possible to match with directories without retaining time data in a database for searching, and time ranges do not need to be specified in a query when searching through the log messages, which is an operation that sometimes takes time. The SQLite3 files where logs are saved are created as FTS (Full-Text Search) tables dedicated to full-text searches, enabling high-speed log retrieval. The SearchEngine accesses the FTS formatted SQLite3 files split each minute to improve the performance of parallel searches. SQL search queries are executed in parallel for each SQLite3 file using GNU Parallel[*4] and results are aggregated via the UNIX pipeline using the awk or count commands.

Hayabusa operates in a stand-alone environment, but it features higher full-text search performance than small-scale Apache Spark[*5] clusters. As shown in Figure 3, a preceding study[*6] indicated that search performance was 27 times faster than a 3-node Apache Spark cluster.

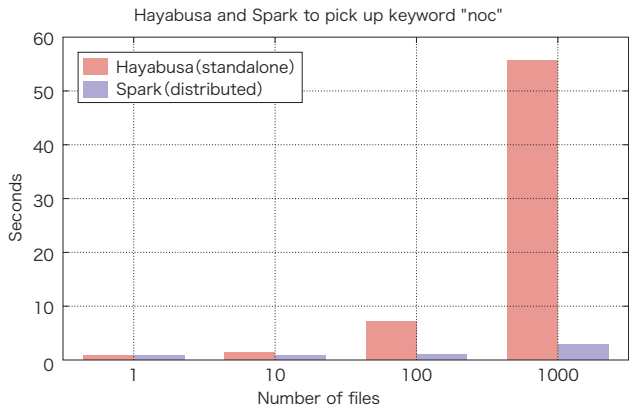

Figure 3: Comparison of Hayabusa and Apache Spark Performance

*3 SQLite3 (https://www.sqlite.org/).
*4 GNU Parallel (https://www.gnu.org/software/parallel/).
*5 Apache Spark (https://spark.apache.org/).
*6 H. Abe, K. Shima, Y. Sekiya, D. Miyamoto, T. Ishihara, and K. Okada. Hayabusa: Simple and fast full-text search engine for massive system log data. In Proceedings of the 12th International Conference on Future Internet Technologies, CFI'17, pages 2:1-2:7, New York, NY, USA, 2017. ACM.

3.4 Hayabusa Distributed Processing

However, stand-alone environments have hardware performance limitations, and it is estimated that performance will be overtaken by other larger scale distributed processing cluster sytems sooner or later. In light of this, with Hayabusa we aimed to remove the constraints of a stand-alone environment, and implemented an architecture capable of scaling out search performance by constructing a distributed processing environment for Hayabusa using multiple hosts.

We redefined the stand-alone version of Hayabusa as a distributed processing system and conducted tests for scaling out search processing capability. To leverage the stand-alone processing capability of Hayabusa, each processing request is sent from the client to the target host as a high-speed RPC (Remote Procedure Call). Parallel searches are implemented on the processing host using GNU Parallel, and results are returned to the client using RPC, and then aggregated. This enables the fast searches in stand-alone environments that Hayabusa originally made possible, while implementing high-speed requests and responses via RPC. To implement a search function that scales out in a distributed host environment, we designed the parallel data storage and distributed search functions separately.

3.4.1 Parallel Storage and Distributed Search

To evaluate the scale out performance of search processes, we configured all the processing hosts to retain the same data, so that searches were possible regardless of which processing host would receive the processing requests. That means syslog data is copied and sent to all of the processing hosts. This ensures that the same result is returned no matter which processing host handles the processing request.

We used RPC to implement search requests to a set of distributed processing hosts. Distributed processing for Hayabusa assumes that processing hosts that receive requests from clients will issue search queries to multiple database files saved to a local disk, much like a stand-alone Hayabusa implementation. It is possible to pass SQLite3 commands to a processing host as a search request parameter, but to achieve a simple distributed process while making full use of the performance of a stand-alone Hayabusa, similar GNU Parallel commands are passed as request parameters in this proposal. In a distributed search, processing requests are issued to multiple processing hosts, so we selected a producer-consumer model where clients queue processing requests, and workers running on processing hosts acquire queued processing requests and return the results. Workers running on each processing host execute the search process after acquiring a processing request.

In this proposal, data is copied to all processing hosts, so the same results are returned no matter which processing host receives a processing request. We also performed load balancing so that processing requests from clients are not concentrated on certain processing hosts, meaning requests are distributed evenly.

3.4.2 Implementation

■ Parallel Storage

In our research, we used the open-source software UDP Samplicator*7 to copy syslog data to all the processing nodes. Figure 4 shows an illustration of the syslog replication process. UDP Samplicator can forward incoming UDP packets to a designated target host without changing the source address. This technique makes it possible for the destination host to receive UDP packets as if the data has been received directly from the source.

Normally, UDP Samplicator performs UDP forwarding using a single process. However, the process load increases when a large volume of syslog data is received, sometimes causing CPU usage to hit 100%. This load can cause packet forwarding to be delayed, leading to dropped packets. In these tests, we patched the UDP Samplicator source code so that it can be run as multiple

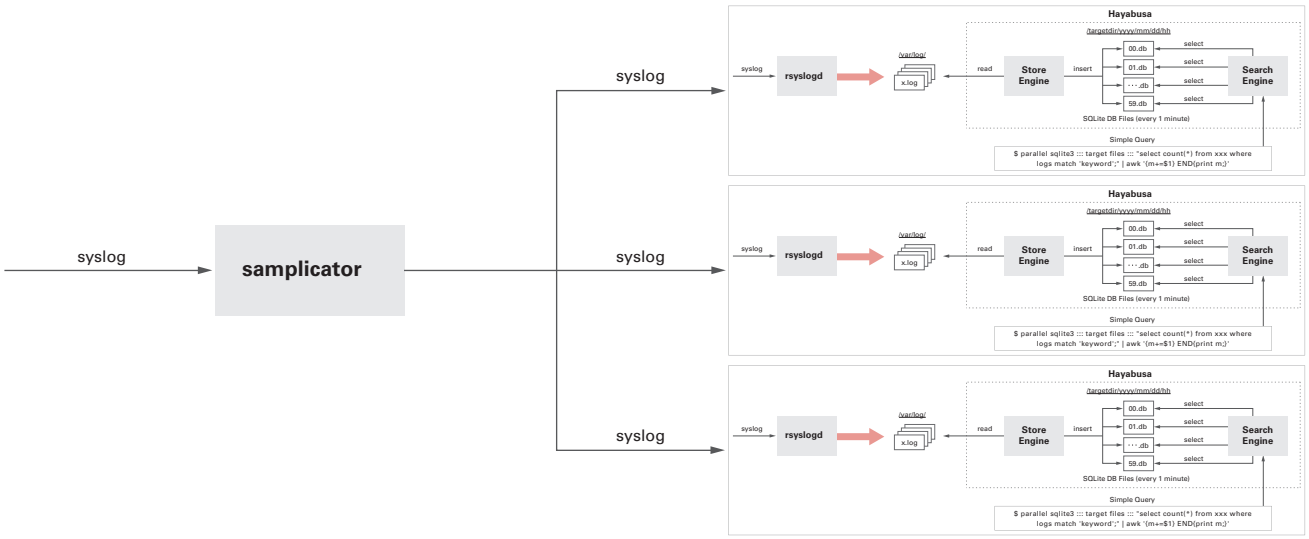

Figure 4: Syslog Replication Using UDP Samplicator

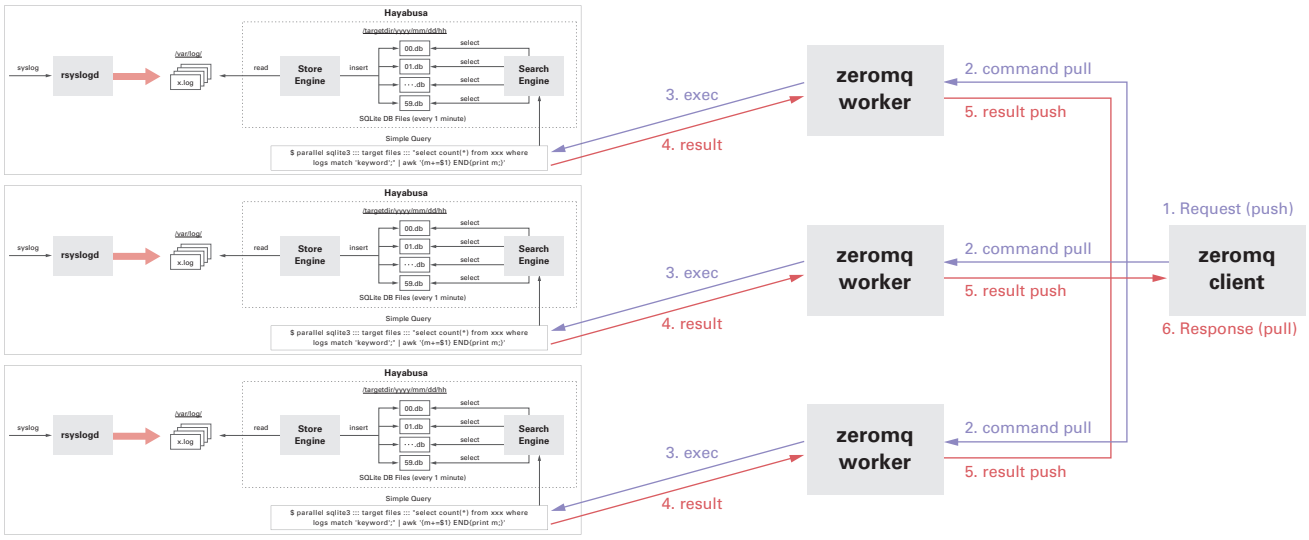

Figure 5: ZeroMQ Push/Pull Implementation Used With UDP Hayabusa

*7 UDP Samplicator (https://github.com/sleinen/samplicator).

processes. More specifically, multiple processes can use the same listening port by adding “SO_REUSEPORT” as a socket option. In this proposal, the incoming syslog port UDP 514 is shared, with packets sent to this port being automatically load balanced between multiple UDP Samplicator processes. This technique makes it possible to scale the replication and forwarding of syslog packets across CPU cores when a large volume of syslog data is received, which tends to become a bottleneck when using only a single CPU core.

■ Distributed Search

The producer-consumer model can be implemented using a wide range of software, but for this research, we used ZeroMQ[*8] because it enables high-speed RPC processing, and client and worker processes can be implemented using libraries. ZeroMQ is used as a distributed message queue that runs at high speed, and it is easy to implement a wide range of messaging patterns such as Request/Response, Publish/Subscribe, and Push/Pull. In this proposal, we implemented a producer-consumer model using the Push/Pull pattern.

As shown in Figure 5, the clients in this proposal were implemented to fulfill both the Push and Pull roles. This enables the client program to perform everything from the issuing of requests, queueing, and the acquisition and aggregation of results through a single client process.

3.5 Evaluation

To investigate the scale-out performance of Hayabusa as a distributed system, we tested whether processing time could be reduced when the number of processing hosts is increased. While increasing the number of processing hosts incrementally from one to ten, the client repeatedly executes 100 requests against one day of data. The size of the data requested 100 times is 14.4 billion records (assuming a syslog flow rate of 100,000 records per minute).

As shown in Figure 6, search processing time for one host was about 468 seconds. As the number of hosts increased, the result was roughly this amount of time divided by the number of hosts. The processing time of ten hosts was 48 seconds.

While 30 million records can be scanned per second when a single host is used, with ten hosts this increases to 300 million, ten times that value. This experiment demonstrates that when the number of hosts is increased from one to ten, search processing performance scales out linearly with the number of hosts.

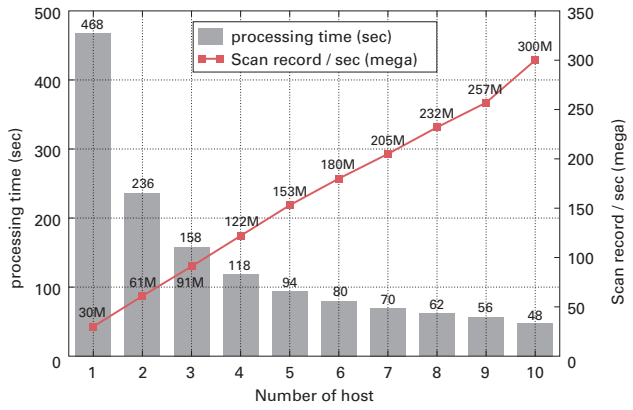

Figure 6: Search Host Scale-Out Performance

*8 ZeroMQ (http://zeromq.org/).

Next, we increased the number of workers incrementally from one process through 16 processes, while processing on ten hosts. We settled on the number 16 because it was the number of physical cores on the CPU of the server used for this test, and we wanted to check whether it would be possible to scale out to the maximum number of physical cores. Running a single worker process on ten hosts took 48 seconds, while running ten worker processes on ten hosts attained a peak processing performance of 6.1 seconds, providing around eight times the processing speed. Processing speed saturated after adding the tenth worker, so we believe attaining maximum performance with sixteen processes may require some additional effort.

When combining the scale-out of both the number of hosts and the number of workers, a search that took 468 seconds when using one worker process on one processing host, dropped to 6.1 seconds when using ten worker processes on ten hosts, which is about 78 times faster.

3.6 Future Challenges

In this research, we copied the same syslog data to each processing host to perform distributed queries in an attempt to improve search performance. This involves the mass replication of data, meaning that as the amount of data increases, there is a waste of network bandwidth and the data to be retained. It is possible to define the number of times that data is replicated, then distribute and retain data across multiple hosts as with HDFS in Hadoop*9. However, in this case, a metadata management system manages the data, and data access is provided through this system as well, which may decrease performance when accessing storage, and lower the performance of the search process. Considering the advantages and disadvantages of Hayabusa and distributed storage, and implementing a distributed storage solution suitable for Hayabusa at the same time, will be a major challenge.

Because Hayabusa operates as a search infrastructure system, it has the potential to provide more benefits when combined with specific application software that runs on it.

It enables high-speed syslog searches, so we believe it will be possible to create high-speed anomaly detection applications when combining with prior research on anomaly detection*10 using syslog data in event networks.

*9 Apache Hadoop (http://hadoop.apache.org/).

*10 Hiroshi Abe and Mikifumi Shikida. Proposal of the anomaly detection method analyzing syslog data using Bollinger Bands algorithm on event network. In Proceedings of the Internet and Operation Technology Symposium 2016, Volume 2016, Pages 57-64, Dec 2016.

3.7 Hayabusa Applications

One of the goals of research in the security field is to predict and detect attacks.

To counter the increasing threat of cyber attacks, it would help to be able to analyze signs of attacks in real time and assess potential attacks as well as their severity and the extent of the impact. This would make it possible to utilize machine learning and deep learning to support responses to security incidents that have up until now been largely dependent on the individual skills of the staff in charge.

The IIJ Research Laboratory is taking part in the "Real-time Attack Detection and Prediction via Cyber Threat Big Data Analysis"*11 project in collaboration with the University of Tokyo, Tokyo Institute of Technology, and Nara Institute of Science and Technology. In this project, Hayabusa will be incorporated as part of the infrastructure for analyzing signs of attacks in real time. It will also operate as infrastructure for gathering statistics and applying machine learning analysis to signs of attack in collaboration with other data analysis and machine learning software (R, Chainer, Pandas, etc.).

3.8 Conclusion

In this report, we provided an overview of the Hayabusa open-source software and experiment results that used distributed processes. The fact that we were able to perform a full scan across 14.4 billion records in approximately six seconds means that we have achieved a full data scan speed on the same level as Google's BigQuery. The ability to realize such a high-speed scan with ten processing hosts demonstrates this is also a reasonable and high-performance distributed processing system from a cost perspective.

A more detailed explanation of the contents in this report have been published in the paper*12.

Part of this research has been carried out with the support of the "Strategic Creation Research Promotion Project (CREST) JPMJCR1783" research and development project implemented by the Japan Science and Technology Agency (JST).

Author:
Hiroshi Abe
Researcher, IIJ Research Laboratory

Author:
Keiichi Shima
Senior Researcher, IIJ Research Laboratory

*11 "Real-time Attack Detection and Prediction via Cyber Threat Big Data Analysis" (https://www.jst.go.jp/kisoken/crest/project/1111094/1111094_13.html) (in Japanese).
*12 Hiroshi Abe and Yoichi Shinoda. Research of the scalable syslog search engine and evaluation. In Proceedings of the Internet and Operation Technology Symposium 2017, Volume 2017, Pages 73-80, Nov 2017.

Internet Topics

JANOG 41 Meeting - The First Hosted by IIJ

The JANOG 41 Meeting held in Hiroshima between January 24 and January 26, 2018, was the first JANOG Meeting*1 hosted by IIJ. Although there was a light dusting of snow during the event, the weather was generally fine. At final count there were 1,171 participants at the main session, and 725 at the social function. These were the highest attendance numbers since the JANOG Meeting was first held, so the event ended in a great success. Here, we will discuss IIJ's initiatives as host of the JANOG 41 Meeting.

Our Role as Host

The JANOG Meeting is a conference held by JANOG (JApan Network Operators' Group), an organization made up of engineers who manage the Internet in Japan. The JANOG Meeting where engineers gather twice a year is organized by an executive committee selected from among JANOG members, but host companies take turns at arranging the venue and coordinating other related matters. The host company goes to great lengths to prepare a venue with suitable amenities to entertain the attendees that come together from all over the country. With IIJ acting as host this time, we focused on the network facilities available at the venue. Because the JANOG Meeting is a hands-on conference concerning Internet-related matters, we naturally use the Internet during the event. The engineers gathered at the venue are also active personnel who handle operations. Although some engineers stay back at the office to hold the fort, at times work is carried out at the venue on a temporary basis. To meet these needs a network environment is required, but there are significant obstacles to implementing this.

High Density and Short Construction Period

Each year the number of JANOG Meeting attendees has been rising, and for the past few years there have consistently been over 600 participants. The event that IIJ hosted had over 1,000 people in attendance. Almost all engineers gathered at the venue use a laptop computer, and in many cases an attendee will use multiple devices such as a smartphone or tablet. It is rare to have such a high density of users gathered in the confined space of a conference hall.

The amount of time that can be spent preparing the network is also quite limited in comparison to the network scale. Because usage fees are expensive for halls that can accommodate 1,000 people, they are only rented for the three days that the JANOG Meeting will be held. That means there was the major constraint that only the morning of the first day of the event could be set aside for preparing the venue network. It is also necessary to pack up equipment promptly after the conference ends.

Providing Internet Access at the Venue

Amidst these constraints, IIJ aimed to provide a pure network experience with high performance.

IIJ's backbone network extends throughout Japan, and we have an NOC in Hiroshima that serves as a base for this. In coordination with Energia Communications, a local provider and fellow JANOG participant, we brought a fiber-optic line that connects directly to the IIJ Hiroshima NOC (Figure 1) into the International Conference Center Hiroshima venue. This fiber connection operates with 10 Gbps of bandwidth using a WDM system installed by IIJ.

We used an IPv4/IPv6 dual stack for the venue network. In coordination with JPNIC, which manages network resources in Japan, IPv4 addresses were configured to distribute global addresses to each device. It is common to use private addresses within a LAN to cut down on IPv4 addresses and ensure network security, but we took this approach because we intended to take on the challenge of constructing a large-scale LAN based on IPv4 global addresses.

Wireless LAN Equipment at the Venue

Because the many attendees take turns using the network, it is essential to construct a wireless LAN network at the venue. In this case we used many SA-W2 wireless LAN-compatible devices developed independently by IIJ. The SA-W2 is a device

*1 JANOG: Stands for JApan Network Operators' Group. This is a group aimed at contributing to Internet engineers and users in Japan by discussing, evaluating, and introducing Internet-based technology and related operational matters.

that can be integrated with the SACM (Service Adapter Control Manager) device management system that IIJ develops and operates. It incorporates mechanisms for reducing the operational burden, such as sending operating status notifications to a management server automatically by simply connecting network and power cables. We used this equipment with the goal of having network construction completed in a short time. We also collaborated with the local Hiroshima City University to provide wireless LAN for the conference. The university aims to improve the efficiency of wireless LAN network utilization, and is researching techniques for estimating the communication quality of wireless access points. For this reason, we set up a device for measuring wireless LAN at the JANOG 41 meeting in the main hall and collected information. For this collection task we use an SA-W2 with special firmware implemented. IIJ has provided wireless LAN at several events and evaluated the service in the past. However, by conducting joint research with Hiroshima City University this time, we expect to gain new insight into techniques for estimating communication quality in wireless LAN networks.

Future Considerations

At the time of writing, the JANOG 41 Meeting has ended, and IIJ has finished providing network services. We are currently evaluating the data gathered through these initiatives, and we will report on the results in the IIJ Engineers Blog*2.

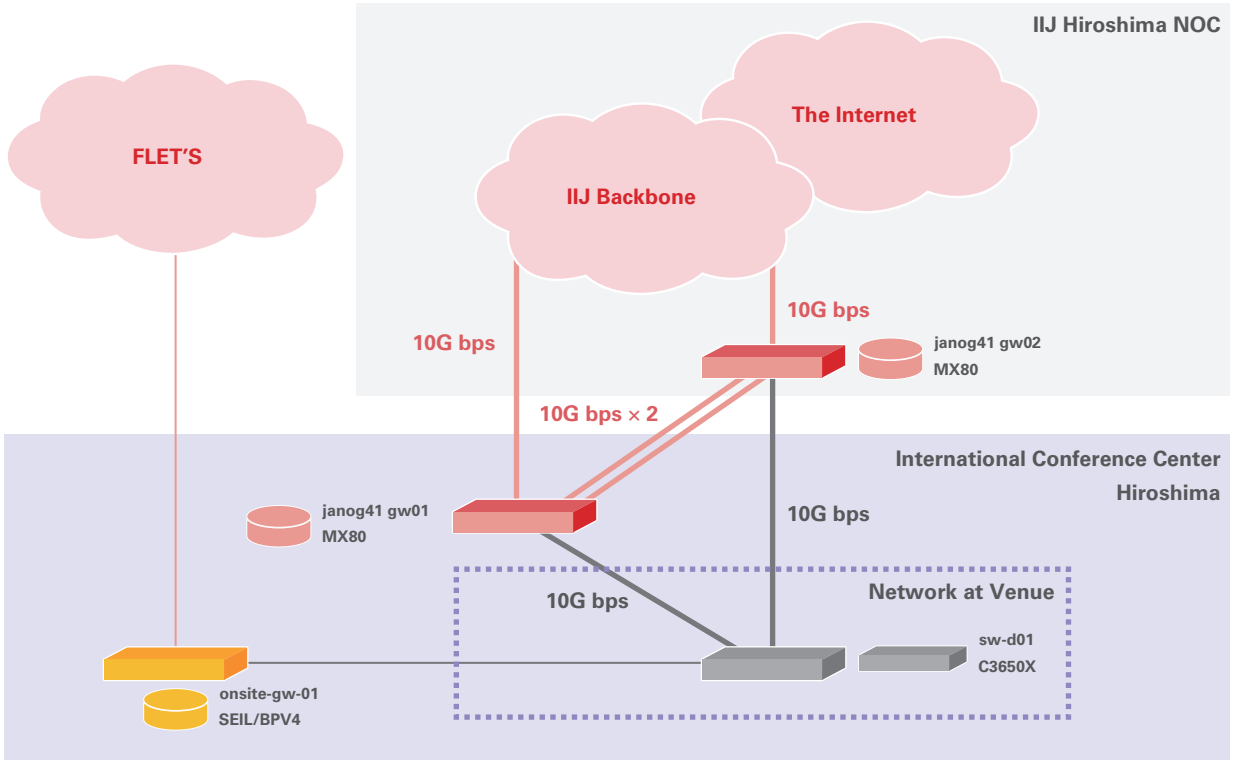

Figure 1: JANOG 41 Meeting Network Configuration

Author:
Kiyotaka Doumae
Manager, Corporate Communications Department, IIJ

*2 IIJ Engineers Blog (http://eng-blog.iij.ad.jp/) (in Japanese). An official blog written by engineers working in the areas of development and operation.

March 2018 Vol.38

IIJ
Internet Initiative Japan

About Internet Initiative Japan Inc. (IIJ)

IIJ was established in 1992, mainly by a group of engineers who had been involved in research and development activities related to the Internet, under the concept of promoting the widespread use of the Internet in Japan.

IIJ currently operates one of the largest Internet backbones in Japan, manages Internet infrastructures, and provides comprehensive high-quality system environments (including Internet access, systems integration, and outsourcing services, etc.) to high-end business users including the government and other public offices and financial institutions.

In addition, IIJ actively shares knowledge accumulated through service development and Internet backbone operation, and is making efforts to expand the Internet used as a social infrastructure.

IIJ-MKTG020-0036

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/