SOC Report

1.1 Introduction

IIJ launched the wizSafe security brand in 2016 and works constantly to create a world in which its customers can use the Internet safely. In our SOC Report in Vol. 38[*1] we examined the Data Analytics Platform at the core of wizSafe, and in Vol. 42[*2] we discussed threats that came to light in 2018 and new initiatives using the Data Analytics Platform. Here, we review key security topics for 2019 (Section 1.2) and discuss observations made on the Data Analytics Platform about threats related to those topics (Section 1.3).

1.2 2019 Security Topics

Key security topics that our SOC focused on in 2019 are summarized in Table 1.

*1 Internet Infrastructure Review (IIR) Vol.38 (https://www.iij.ad.jp/en/dev/iir/038.html).

*2 Internet Infrastructure Review (IIR) Vol.42 (https://www.iij.ad.jp/en/dev/iir/042.html).

Table 1: Key security topics in 2019

Month	Summary
January	**A personal information leak occurred on a file transfer service run by a Japanese internet services company due to unauthorized access by a third party. Roughly 4.8 million rows of member data were affected, and it was announced that the service would close on March 31, 2020.** "Closure of the Taku-File-Bin service (January 14, 2020) " (retrieved January 14, 2020) https://www.filesend.to/ (in Japanese)
February	**In February 2019, Japan's Ministry of Internal Affairs and Communications and the National Institute of Information and Communications Technology (NICT) launched a project called NOTICE (National Operation Towards IoT Clean Environment) to survey IoT devices, find those vulnerable to use in cyberattacks (e.g., due to weak passwords), and alert the users of those devices.** "The 'NOTICE' Project to Survey IoT Devices and to Alert Users" https://www.soumu.go.jp/main_sosiki/joho_tsusin/eng/Releases/Telecommunications/19020101.html "The 'NOTICE' Project to Survey IoT Devices and to Alert Users" https://www.nict.go.jp/en/press/2019/02/01-1.html
March	**The Coinhive service ended on March 8. The reason given was that factors such as repeated changes to cryptocurrency specifications and a decline in market value made it financially difficult to continue the service.**
March	**Servers run by a foreign PC manufacturer were subject to an APT (Advanced Persistent Threat). As a result, files containing malicious code were transmitted to some users who ran updates using the utilities bundled with the manufacturer's notebooks.** "ASUS response to the recent media reports regarding ASUS Live Update tool attack by Advanced Persistent Threat (APT) groups" http://www.asus.com/News/hqfgVUyZ6uyAyJe1
April	**It was discovered that an "ac.jp" domain (reserved for use in Japan by higher education institutions etc.) had been acquired by a non-qualified third party and that the domain had been used to host an adult website. The reported cause was inadequate checking of the registree's eligibility to register the domain. Given the need to ensure the credibility of highly public domains, the Ministry of Internal Affairs and Communications ordered that steps be taken to prevent a recurrence.** "Administrative action (order) relating to Japan Registry Services Co., Ltd.'s management of '.jp' domain names" https://www.soumu.go.jp/menu_news/s-news/01kiban04_02000152.html (in Japanese)
May	**A remote-code execution vulnerability in Remote Desktop Services, commonly known as BlueKeep, was revealed. As this was judged to have a serious impact on the spread of malware, an update was provided for end-of-life OS versions. Attacks actually using BlueKeep were also observed in November.** "CVE-2019-0708 | Remote Desktop Services Remote Code Execution Vulnerability" https://portal.msrc.microsoft.com/en-US/security-guidance/advisory/CVE-2019-0708
May	**It was announced that three IT security companies had been hacked and that confidential information including development documentation and antivirus source code may have been stolen. It was later concluded that one of the companies had not been impacted by the incident.** "Top-Tier Russian Hacking Collective Claims Breaches of Three Major Anti-Virus Companies" https://www.advanced-intel.com/post/top-tier-russian-hacking-collective-claims-breaches-of-three-major-anti-virus-companies
June	**A number of FreeBSD and Linux kernel vulnerabilities related to TCP were announced, including the vulnerability commonly known as SACK Panic (CVE-2019-11477), which could allow a kernel panic to be triggered by the receipt of deliberately crafted SACK packets.**
July	**It was announced that some accounts on a barcode-based payment service had been subject to unauthorized access and use by third parties. The reason given was inadequate restrictions against logging in on multiple devices and insufficient additional authentication, including two-step authentication. The service was terminated on September 30 in response.** "Notice of 7pay service termination, background, and response going forward" https://www.sej.co.jp/company/important/201908011502.html (in Japanese)
July	**It was announced that around 3 billion yen worth of cryptocurrency had been taken from a Japan-based cryptocurrency exchange. The funds taken were stored in "hot wallets", which are kept in online environments, and it is thought that the private keys had been stolen and used without authorization.** "(Update) Notification and Apology Regarding the Illicit Transfer of Crypto Currency at a Subsidiary of the Company (Third Report)" https://contents.xj-storage.jp/xcontents/AS08938/0bf3e2e9/7a8a/4e9f/97d5/0f0a146233de/20190802124804913s.pdf
July	**It was disclosed that an Elasticsearch (full-text search engine) database containing a Japanese automaker's internal information had been left open to unauthenticated access. The roughly 40GB of information included employees' personal information as well as information on the internal network and devices.** "Honda Motor Company leaks database with 134 million rows of employee computer data" https://rainbowtabl.es/2019/07/31/honda-motor-company-leak/
August	**Increase in attacks targeting a vulnerability in several SSL VPN products announced in April 2019 onward. Details on the vulnerability were revealed at Black Hat USA 2019 in August, and observations of PoC exploits and attacks using this vulnerability were also reported. Our SOC also observed attack traffic exploiting the Pulse Secure vulnerability (CVE-2019-11510).** "Over 14,500 Pulse Secure VPN Endpoints Vulnerable to CVE-2019-11510" https://badpackets.net/over-14500-pulse-secure-vpn-endpoints-vulnerable-to-cve-2019-11510/
September	**It was reported that an Elasticsearch (full-text search engine) database containing the information of over 20 million Ecuadorians had been left open to unauthenticated access.** "Report: Ecuadorian Breach Reveals Sensitive Personal Data" http://www.vpnmentor.com/blog/report-ecuador-leak/
September	**DDoS attacks were launched on Wikipedia, Twitch, and Blizzard servers. The attacks were staged by a botnet thought to be a Mirai variant.**
November	**JPCERT/CC issued an alert on the Emotet malware. The organization said that it had received multiple reports from late October 2019 of infections caused by Word files attached to forged emails purporting to be from actual organizations or people. And our SOC observed increased levels of such activity from end-September 2019.** "Alert Regarding Emotet Malware Infections" https://www.jpcert.or.jp/english/at/2019/at190044.html
December	**It was announced that several companies had been infected by the Emotet malware. Alerts were sent out saying that email addresses and email text saved on the infected devices may have been leaked and that people should not open attachments or URLs in suspicious emails purporting to be from any of the companies affected.**
December	**It was discovered that hard disks had been stolen from leased servers returned by a local government at the end of the lease before the data had been deleted from them. The hard disks were taken by an employee of the company hired by the leasing firm to erase the data and auctioned off on an online auction site.** "Theft of harddisks returned after lease expiry" http://www.pref.kanagawa.jp/docs/fz7/prs/r0273317.html (in Japanese)
December	**A report indicated that over 267 million user records on a foreign social networking service were left exposed on an Elasticsearch server that was publicly accessible without authentication.** "Report: 267 million Facebook users IDs and phone numbers exposed online" http://www.comparitech.com/blog/information-security/267-million-phone-numbers-exposed-online/

1.3 Observational Data

This section looks at notable activity in 2019 as revealed using the Data Analytics Platform.

1.3.1 Information Leaksfrom Externally Exposed Elasticsearch Servers

■ Elasticsearch and Information Leaks

Large-scale breaches of personal information were frequent in 2019. Particularly notable were information leaks due to poorly configured Elasticsearch (full-text search engine) servers. The security topics in Section 1.2 included three information leaks related to Elasticsearch. In addition to the cases listed there, an Elasticsearch server containing a U.S.-based cloud data management company's customer information was left externally accessible without authentication, according to a report[*3] in January 2019, and likewise for an Elasticsearch server containing information on roughly 90% of Panama citizens, per a May 2019 report[*4]. Large amounts of information were leaked in both cases, with the number of records exceeding several million and the volume of data exceeding several dozen GB.

Elasticsearch is an open-source, full-text search engine based on Apache Lucene developed primarily by Elastic[*5]. It employs parallel processing of massive datasets on distributed systems to achieve its high-speed search capabilities. This makes it a not-uncommon choice for large information banks, and is also why large amounts of information tend to be leaked when security incidents do occur. Elasticsearch provides a RESTful API and allows searches and data operations to be performed over the HTTP protocol. The default port for HTTP access is 9200/TCP.

In 2019, our SOC observed an increase in scanning traffic on port 9200/TCP that we believe indicates searches for Elasticsearch servers.

■ Observational Data

Figure 1 plots the number of 9200/TCP scans and source IP addresses observed over time on the IIJ Managed Firewall Service. The number of scans is normalized to a percentage of total 9200/TCP scans observed over the full year such that the overall total is 100%.

Figure 1 shows a noticeable rise in scans over September 21 – October 31. Scans of 9200/TCP over these 41 days accounted for roughly 23.37% of all scans during 2019, and the number of source IP addresses per day rose to a peak of 30,394. This is about 98.68 times the IP address count for January 1 (308). No Elasticsearch vulnerabilities

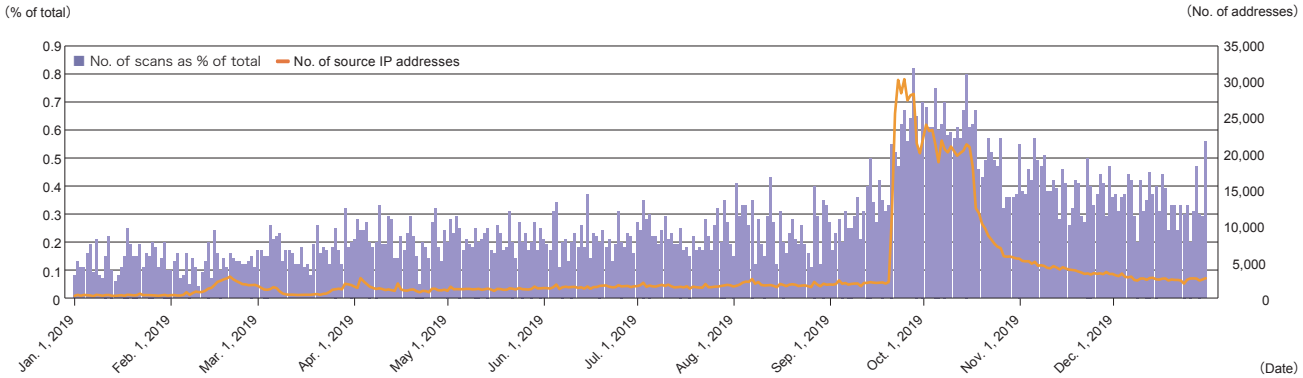

Figure 1: Scanning of 9200/TCP (January–December 2019)

*3 TechCrunch (https://techcrunch.com/2019/01/29/rubrik-data-leak/).

*4 Security Affairs, "Personally identifiable information belonging to roughly 90% of Panama citizens were exposed on a poorly configured Elasticsearch server" (https://securityaffairs.co/wordpress/85462/data-breach/panama-citizens-massive-data-leak.html).

*5 Elasticsearch (https://www.elastic.co/).

were announced around this time, and we have found no clear reason for the increased scanning activity.

On September 16, just before the spike in scans, it was reported that information on more than 20 million Ecuadorians had been exposed. While it is possible that this is what prompted the spike, we can find no clear evidence linking the two phenomena.

Temporary increases in scanning activity also occurred in mid-February and early April. The source IP count briefly rose to nearly 3,000 on February 20 and April 3. On February 19, Elastic announced an Elasticsearch vulnerability (CVE-2019-7611)[*6], and we surmise that the upticks represent searches for servers running Elasticsearch. CVE-2019-7611 is an access permission vulnerability that, if exploited, could allow information acquisition or tampering. Further, Talos, Cisco Systems' security arm, published a report on Elasticsearch attacks that occurred around this time[*7]. According to the report, attacks targeting previously revealed vulnerabilities (CVE-2014-3120, CVE-2015-1427) were observed.

The trend across the year also shows an overall rise in 9200/TCP scanning. The average number of scans per day in December rose to 0.35% from the January average of 0.14%, a roughly 2.46-fold rise, and the source IP address count increased roughly 7.02-fold, from 384.74 on average in January to 2702.10 in December. A report on observational data issued by Japan's National Police Agency (NPA)[*8] shows a similar trend and, like the Talos report, discusses attacks thought to have been targeted at CVE-2015-1427.

■ Countermeasures

Most of the widely reported information leaks involving Elasticsearch in 2019 were due to poorly configured servers allowing authentication-free access. Important basic countermeasures include using a firewall to exclude unnecessary traffic, including on 9200/TCP, if the system does not need to be accessible from the Internet and setting up appropriate authentication to only permit connections from trusted IP addresses. And as the Talos and NPA reports indicate, attacks targeting past vulnerabilities continue to be observed. When information on vulnerabilities relevant to your system is released, you need to determine what the impact on your system is and apply the relevant patches.

1.3.2 DDoS Attack Observations

IIJ observes and responds to DDoS attacks employing various methodologies. This section summarizes key topics in DDoS attacks in 2019. We start by looking at attacks detected by the IIJ DDoS Protection Service in 2019. Next, we look at attack methods that were much talked about in 2019. And finally we go over examples of damage caused by those attack methods in 2019, along with observational data.

*6 Elastic, "Security issues" (https://www.elastic.co/jp/community/security).

*7 Cisco Talos, "Cisco Talos Honeypot Analysis Reveals Rise in Attacks on Elasticsearch Clusters" (https://blog.talosintelligence.com/2019/02/cisco-talos-honeypot-analysis-reveals.html).

*8 National Police Agency, "Increase in online traffic aimed at Elasticsearch vulnerability" (in Japanese, https://www.npa.go.jp/cyberpolice/important/2019/201910021.html).

Summary of 2019 DDoS Attack Observations

DDoS attacks on Wikipedia, Twitch, and Blizzard created a stir in September 2019. Of the DDoS attacks IIJ responded to in 2019, here we summarize those detected by the IIJ DDoS Protection Service. Table 2 shows the number of attacks and traffic volume detected by the IIJ DDoS Protection Service.

Of the attacks in Table 2, the SYN Flood and SYN/ACK attacks use TCP, and the UDP Amplification and UDP Flood attacks use UDP. A number of application protocols are used in UDP Amplification attacks, including DNS, NTP, and LDAP.

Table 2 shows the daily average number of attacks for each month. No month in 2019 was a particular standout for DDoS attacks. May recorded the highest number of packets per second, and the longest attack occurred in January. The maximum number of packets was relatively large in May, July, and December, but the longest attacks in those months were under one hour. UDP Amplification attacks using LDAP and DNS feature prominently in the maximum traffic and maximum attack duration listings.

Key DDoS Attack Topics for 2019

A number of new methodologies suited to DDoS attacks other than those appearing in Table 2 also popped up in 2019. Three keywords stood out on the DDoS landscape in 2019.

- Web Services Dynamic Discovery (WSD)
- Apple Remote Management Service (ARMS)
- SYN/ACK reflection

The first, WSD, is a protocol that uses the Simple Object Access Protocol (SOAP) to locate services and enable data exchanges in specific network ranges. It uses port 3702/UDP, and it is known to be used on printers and PCs that run on Windows Vista and up. The possibility of DDoS attacks using this protocol has been discussed by zeroBS GmbH[*9]. It has been observed that there are roughly 630,000 IP addresses online that respond on 3702/UDP[*10]. Our SOC observed an increase in 3702/UDP scanning activity in August 2019[*11]. Figure 2 shows scanning activity on this port observed at the SOC in 2019. Note that the

Table 2: Summary of Observational Data on DDoS in 2019

Month	No. of incidents (daily avg.)	Approx. max. packets/sec. (x10,000)	Maximum traffic		Maximum attack duration	
			Bandwidth	Method	Duration (h:mm)	Method
1	13.58	~179	17.38Gbps	DNS Amplification	3:20	SYN Flood
2	15.75	~284	27.89Gbps	LDAP Amplification	1:18	LDAP Amplification
3	14.00	~652	19.30Gbps	SSDP Amplification	2:32	SSDP Amplification
4	22.96	~97	9.21Gbps	DNS Amplification	0:41	DNS Amplification
5	16.16	~886	39.29Gbps	LDAP Amplification	0:41	DNS Amplification
6	10.93	~148	8.11Gbps	SSDP Amplification & SYN/ACK reflection	0:30	SSDP Amplification & SYN/ACK reflection
7	16.41	~738	75.67Gbps	DNS Amplification	0:38	NTP Amplification
8	18.10	~91	8.77Gbps	LDAP & DNS Amplification	1:35	UDP Flood
9	19.20	~130	11.71Gbps	LDAP & DNS Amplification	0:43	NTP Amplification
10	22.09	~310	23.09Gbps	Amplification: LDAP, DNS, NTP, etc.	1:56	LDAP Amplification
11	13.36	~70	8.24Gbps	UDP Flood	0:25	UDP Flood
12	10.38	~607	61.34Gbps	LDAP & DNS Amplification	0:38	NTP Amplification

*9 zeroBS, "Analysing the DDOS-Threat-Landscape, Part 1: UDP Amplification/Reflection" (https://zero.bs/analysing-the-ddos-threat-landscape-part-1-udp-amplificationreflection.html).

*10 zeroBS, "New DDoS Attack-Vector via WS-Discovery/SOAPoverUDP, Port 3702" (https://zero.bs/new-ddos-attack-vector-via-ws-discoverysoapoverudp-port-3702.html).

*11 wizSafe, "wizSafe Security Signal August 2019 Observational Report" (in Japanese: https://wizsafe.iij.ad.jp/2019/09/746/).

number of scans is normalized to a percentage of total 3702/UDP scans observed over the full year such that the total is 100%.

Figure 2 shows that scans on this port increased from around August 13. The number of source IP addresses scanning the port also rose from August 19 through end-August. The observations in Figure 2 generally match those in BinaryEdge reports. The reason for the increase in scanning on the port on February 17 is unclear, but Baidu, Inc. reported on February 19 that a DDoS attack using WS-Discovery had occurred[*12]. Hence, it appears that DDoS attacks exploiting WS-Discovery had been in use since at least February. But it was September 2019 when they came into focus in Japan. And a US-Cert document on UDP Amplification Factors was updated in December to cite a September article on this type of attack[*13]. So it seems that it was actually a few months after WS-Discovery was first used in attacks that attackers started to use the protocol for DDoS attacks in earnest.

The second keyword, ARMS, is a service used on Apple Remote Desktop (ARD). ARD is an application for remotely controlling macOS devices. ARMS receives commands from the control console via 3283/UDP. It was found that there are around 40,000 devices on which ARMS is reachable via the Internet[*14]. Figure 3 shows scanning activity on the port observed by our SOC in 2019. Note that the number of scans is normalized to a percentage of total 3283/UDP scans observed over the full year such that the total is 100%.

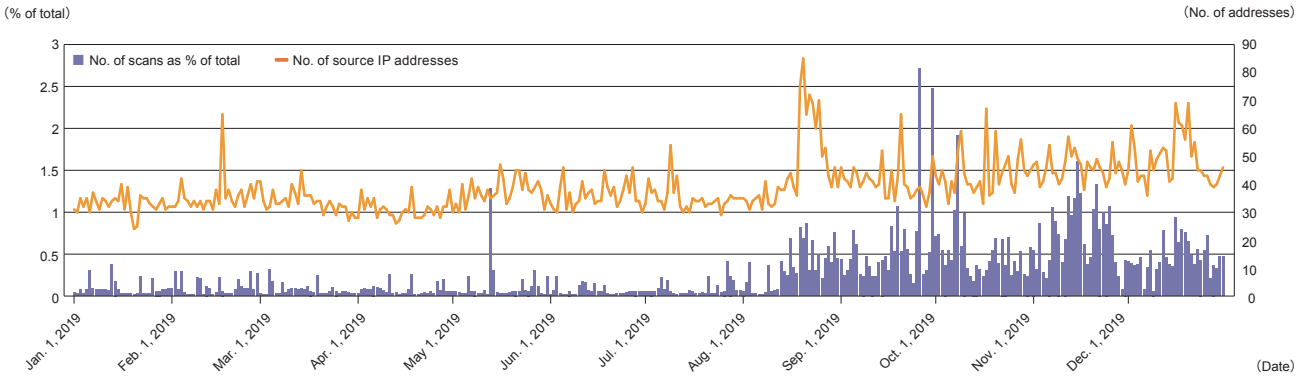

Figure 2: Scanning of 3702/UDP and Number of Source IP Addresses (Jan.–Dec. 2019)

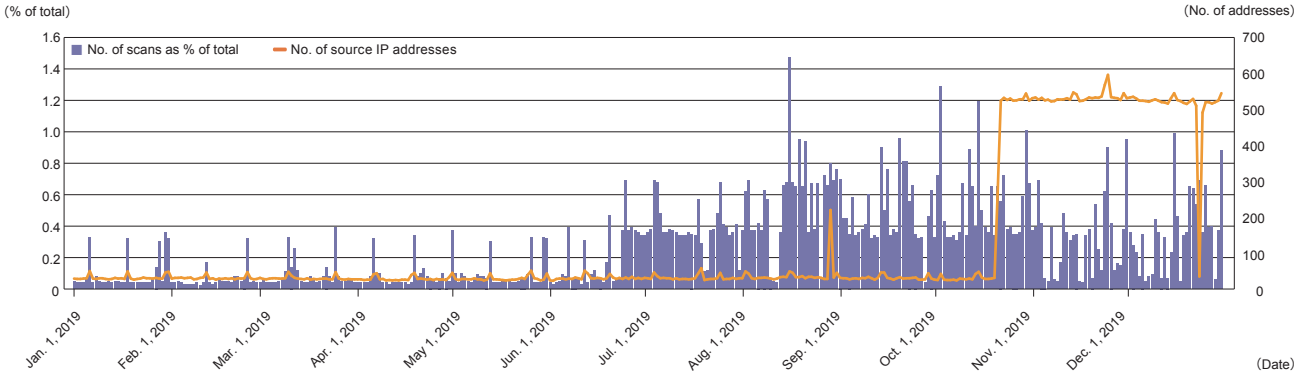

Figure 3: Scanning of 3283/UDP and Number of Source IP Addresses (Jan.–Dec. 2019)

*12 Baidu Security Index,"基于ONVIF协议的物联网设备参与DDoS反射攻击"(in Chinese, https://bsi.baidu.com/article/detail/128).
*13 CISA, "Alert (TA14-017A)" (https://www.us-cert.gov/ncas/alerts/TA14-017A).
*14 ZDNet, "macOS systems abused in DDoS attacks" (https://www.zdnet.com/article/macos-systems-abused-in-ddos-attacks/).

Figure 3 indicates that scans of the port increased from around June 24. And the number of source IP addresses scanning the port increased from around October 22. So it is evident that scanning activity was increasing a few days before the release of the NetScout Systems, Inc. report[*15].

Our third keyword is SYN/ACK reflection attacks. This attack takes place in the TCP three-way handshake. SYN packets with a spoofed source address are sent to many addresses simultaneously, thereby effectively recruiting the resulting SYN/ACK packet responses to perform a DDoS attack on the source address. Figure 4 gives an overview of a SYN/ACK reflection attack.

Below, we describe the flow of events from the launch of a SYN/ACK reflection attack through to the damage it inflicts on the victim. Refer to Figure 4 as you read through.

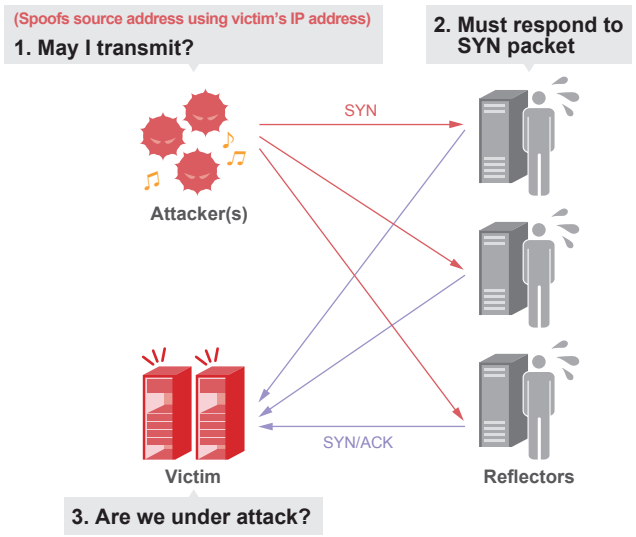

Figure 4: Overview of a SYN/ACK Reflection Attack

1. To generate the SYN/ACK packets used in the attack, the attacker spoofs the source address to match the attack target and sends the SYN packets with that spoofed source address to the reflectors.
2. During the three-way handshake, the reflectors send SYN/ACK packets in response to those SYN packets.
3. Because the source address on the SYN packets is spoofed, the SYN/ACK packet responses from the reflectors are delivered to the attack target's IP address, thus consummating the attack.

This type of attack was observed by our SOC in 2018 and is explained in Section "1.2.2 SYN/ACK Reflection Attack" of Vol. 42[*16]. This SYN/ACK reflection attack uses a TCP Amplification attack technique that was known around 2006[*17]. In 2014, researchers discovered devices on the Internet with protocol implementations that result in more SYN/ACK packets, RST packets, or PSH packets being retransmitted than is common[*18]. It is not clear whether the devices found in 2014 are actually being used, but the attack principles are the same. At our SOC, TCP Amplification attacks that use SYN/ACK packets are termed SYN/ACK reflection attacks, and they were observed frequently from around July through November.

A distinctive feature of these three attack methods that featured prominently in 2019 is that they spoof the packet source address to match the target and recruit reflectors to mount the DDoS attacks. DDoS attacks like this are called Distributed Reflection Denial of Service (DRDoS). To perform a DRDoS attack, the attacker first looks for hosts and ports that can be used as reflectors and attempts to exploit them. So if ports that can be used for DRDoS are made accessible to anyone on the Internet, they are at risk of being recruited as reflectors in DRDoS attacks. With DRDoS attacks like WSD and ARMS, countermeasures are needed not only on the sender and target but also on the reflectors. In DRDoS attacks, the administrators of the reflector servers are not being targeted, but they are unintentionally participating in attacks on the targeted servers or networks. So it is important to make sure you do not unnecessarily leave

*15 NETSCOUT, "A Call to ARMS: Apple Remote Management Service UDP Reflection/Amplification DDoS Attacks" (https://www.netscout.com/blog/asert/call-arms-apple-remote-management-service-udp).
*16 Internet Infrastructure Review (IIR) Vol. 42 (https://www.iij.ad.jp/en/dev/iir/042.html).
*17 RFC 4732, "Internet Denial-of-Service Considerations" (https://tools.ietf.org/html/rfc4732#section-3.1).
*18 USENIX, "Hell of a Handshake: Abusing TCP for Reflective Amplification DDoS Attacks" (https://www.usenix.org/system/files/conference/woot14/woot14-kuhrer.pdf).

ports open on the Internet and configure hosts to allow only the intended access. Not only does this reduce unauthorized access, it also helps reduce the number of reflectors available for DDoS attacks and thus makes it more likely that we can limit attackers' options for staging DDoS attacks. There is a DRDoS attack, however, for which it is not easy to restrict access to reflectors. This is the SYN/ACK reflection attack. The reason for this is explained in the next section on the SOC's observations.

Our SOC's Observations

DDoS attacks using the three methods described targeted organizations and services in Japan in 2019. Here, we look at some of the more prominent DDoS attacks that occurred in Japan in 2019, together with information observed by our SOC. DDoS attacks using WSD and ARMS aimed at organizations in Japan were highlighted in a JPCERT/CC alert in October 2019[*19]. It is known that in these cases, not only were WSD and ARMS used for DDoS attacks but extortion emails demanding cryptocurrency payments were also received. Attempts apparently motivated by monetary gain and involving messages threatening to launch DDoS attacks like this are called Ransom Denial of Service (RDoS) attacks. RDoS attacks caused a stir not only in 2019 but in 2017 as well[*20]. Whether the actors behind the attacks were the same in both years is unclear, but it is at least true that DDoS attacks using WSD and ARMS, which had not been disclosed in 2017, were used in the 2019 cases. Considering this in conjunction with Figures 2 and 3, it appears that DDoS attack infrastructure is being progressively adapted to exploitable protocols.

An example of a SYN/ACK reflection attack being used in a DDoS attack aimed at companies in Japan is that listed for maximum traffic volume and attack length for June in Table 2. A key feature of SYN/ACK reflection attacks is that they use any TCP port as the reflector and thus do not exploit services tied to specific ports like WSD or ARMS. This is why ports commonly used by Web servers, such as 80/TCP and 443/TCP, are used. It is important, for example, that the content of Web servers on the Internet be accessible from anywhere if it is to be made available to a wide audience. In this scenario, the firewall will be configured to allow anyone to access the server. And as such, it will be difficult to deny access on the server side if the server is used as a reflector in a SYN/ACK reflection attack. Figure 5 shows the percentage breakdown of TCP ports used as reflectors in SYN/ACK reflection attacks observed by our SOC in 2019.

As Figure 5 shows, the TCP ports used in SYN/ACK reflection attacks are 80/TCP, 443/TCP, and 25/TCP, which is used for the Simple Mail Transfer Protocol (SMTP). These account for over 95% of the total. The "Others" slice represents many ports including 21/TCP, 22/TCP, and 587/TCP. So it is evident that the ports recruited to stage SYN/ACK reflection attacks are TCP ports that are relatively openly accessible on the Internet. As Figure 5 shows, TCP ports on which services that permit external access are running appear prone to exploitation, making it difficult to deal with SYN/ACK reflection attacks by implementing access controls on the reflectors.

Yet this is not the only challenge in dealing with SYN/ACK reflection attacks. SYN/ACK reflection attacks are tough to identify unless you basically have an overview of the entire network, encompassing all the devices, the attack target, and so on. Since a slew of SYN packets from the attacking device actually arrives at each host recruited as a reflector, the reflector host administrators are liable to conclude that a SYN Flood attack is underway. In that case, if the source address in the SYN packets is permanently blacklisted on reflector hosts, it will not be possible to reach those hosts from the

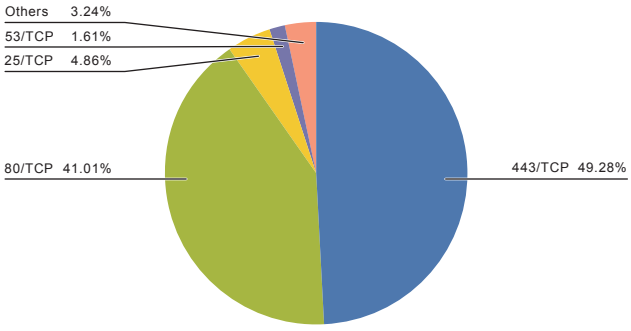

Figure 5: Breakdown of Reflector Ports Used in SYN/ACK Reflection Attacks

*19 JPCERT/CC, "Extortion emails threatening DDoS attacks and demanding cryptocurrency" (in Japanese, https://www.jpcert.or.jp/newsflash/2019103001.html).
*20 Radware, "Fancy Bear DDoS for Ransom" (https://security.radware.com/ddos-threats-attacks/threat-advisories-attack-reports/fancybear/).

attack target's IP address once the DDoS attack is over. This is the collateral damage of SYN/ACK reflection attacks.

Examples of devices in Japan being used as reflectors in SYN/ACK reflection attacks are available on our SOC's reporting site, wizSafe Security Signal[*21][*22][*23][*24]. Note that because these are SYN/ACK reflection attacks observed from the reflector's point of view, not the target's, the information does not indicate the full scale of SYN/ACK reflection attacks.

1.3.3 Emotet

■ Overview of Emotet

A malware program called Emotet, which infects hosts by exploiting emails, came to the fore in the latter half of 2019. This malware was first reported[*25] in 2014 by Joie Salvio, then working at Trend Micro. Emotet was initially active as a banking trojan targeting information from financial institutions but bit by bit morphed into a botnet. It also acquired worm capabilities by adopting a modular framework, giving it the ability to spread various malware and ransomware payloads. It has thus morphed in recent years and gained the ability to download malware (Trickbot, ZeuS, etc.) that steals not only financial institutions' information but other confidential information as well. It has also been reported that malware with information stealing capabilities downloaded by Emotet can infiltrate target systems and eventually deploy a ransomware payload called Ryuk. There have been reports of activity dubbed a triple threat[*26] involving a multistage attack in which information stolen by these malware programs is used to infiltrate target systems, on which a ransomware payload called Ryuk is then deployed. As these changes have unfolded, the range of attack targets has also shifted to public institutions and private companies.

Internationally, it was observed[*27] that C2 servers used by Emotet went inert from June 2019, but the hiatus did not last long. It was reported at the end of August 2019 that the servers had resumed activity, and from September on IIJ's email gateway service, the IIJ Secure MX Service, we detected an increase in malicious emails designed to spread Emotet infections.

Our SOC observed a lot of infection activity exploiting Microsoft Word (doc) format attachments. Subsequently, there was an increase in the number of emails representing a separate infection vector, namely that the body text contained a URL that downloads a doc file that then infects the host with Emotet.

*21 wizSafe, "wizSafe Security Signal July 2019 Observation Report" (in Japanese, https://wizsafe.iij.ad.jp/2019/08/717/).
*22 wizSafe, "Observation of DDoS attacks targeting Servers.com" (in Japanese, https://wizsafe.iij.ad.jp/2019/10/764/).
*23 wizSafe, "Examples of TCP SYN/ACK Reflection Attack Observations for October 2019" (in Japanese, https://wizsafe.iij.ad.jp/2019/12/820/).
*24 wizSafe, "Examples of TCP SYN/ACK Reflection Attack Observations for November 2019" (in Japanese, https://wizsafe.iij.ad.jp/2019/12/839/).
*25 Trend Micro, "New Banking Malware Uses Network Sniffing for Data Theft" (https://blog.trendmicro.com/trendlabs-security-intelligence/new-banking-malware-uses-network-sniffing-for-data-theft/).
*26 Cybereason, "Research by Noa Pinkas, Lior Rochberger, and Matan Zatz" (https://www.cybereason.com/blog/triple-threat-emotet-deploys-trickbot-to-steal-data-spread-ryuk-ransomware).
*27 Bleeping Computer, "Emotet Botnet Is Back, Servers Active Across the World" (https://www.bleepingcomputer.com/news/security/emotet-botnet-is-back-servers-active-across-the-world/).

Opening the Emotet-infecting doc file with Word's default settings produces a message asking you to "Enable Content", as in Figure 6. Enabling this results in a macro being executed. If Word is already configured to enable macros, the user does not see a screen like that in Figure 6 and the macro simply runs automatically. Once executed, the macro downloads Emotet from a malware distribution server, which infects your device.

Once it has infected a device, Emotet tries to make the infection persistent by copying itself to new services, configuring them to run automatically. It then steals information from the infected PC and communicates with its C2 server. The information stolen includes email text and addresses, and some of Emotet's malicious emails exploit this information to disguise themselves as replies to past emails threads. This is one factor behind Emotet's spread. As the multi-stage attack (triple threat) example demonstrates, Emotet serves as an entry point for other malware, so the type of damage it ultimately causes is likely to continue to morph ahead.

Observational Data

Below, we report on our SOC's observations on Emotet.

The stacked bar graph in Figure 7 divides attacks detected over September–December 2019 into those related to Emotet and those related other attacks. Date is on the horizontal axis. The vertical axis represents the total number of detections normalized to a percentage of total detections over the entire period, such that the overall total is 100%.

The first prominent Emotet detection in the graph is on September 27. Following that, it was also detected prominently on October 16, 17, 23, and 24. In November onward, it was detected on more days and more frequently than in the preceding months. And the detections tended to be concentrated on weekdays. Detections reached the overall peak for the period over December 3–4. This was followed by a spike on December 16, and then detections on the IIJ Secure MX Service settled down through the rest of December.

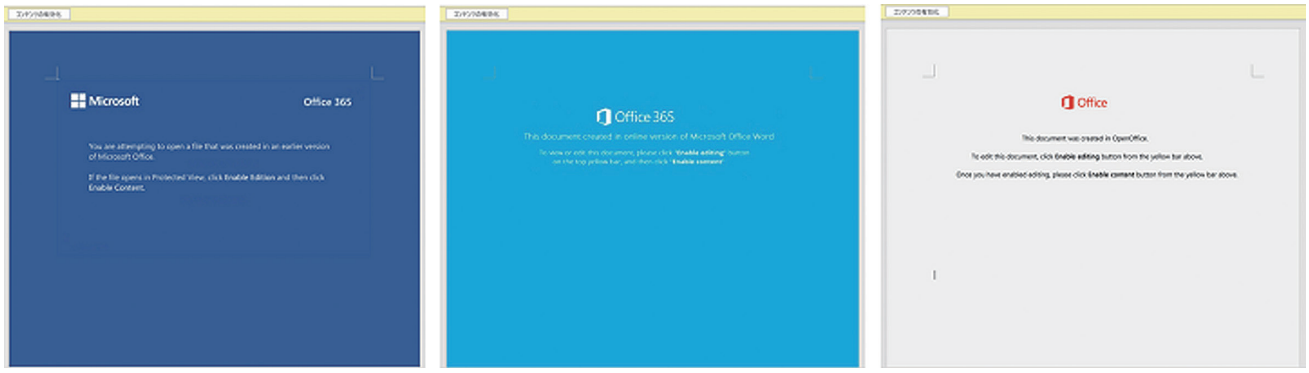

Figure 6: Examples of Screens Asking User to Enable Content

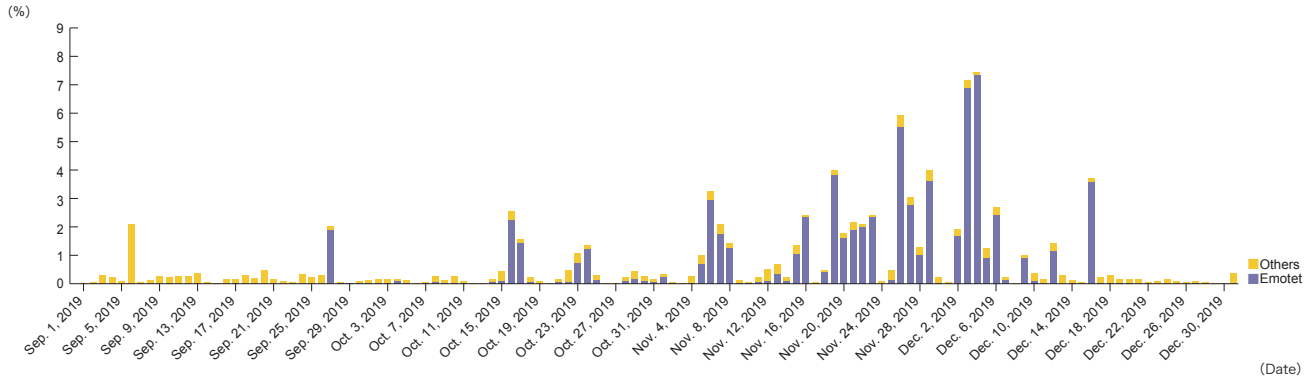

Figure 7: Malware Detections in Received Emails (Sep.–Dec. 2019)

But as if substituting for this, Emotet-related detections on the IIJ Secure Web Gateway Service then increased. In Figure 8, date is on the horizontal axis, and the vertical axis represents the number of detections normalized to a percentage of total in December 2019, such that the total is 100%.

These Emotet-related detections in Figure 8 increased for a few days starting December 17, right after the email detections in Figure 7 eased off. We have determined that this traffic represents attempts to download Emotet-infecting doc files. Japan's Information-technology Promotion Agency also issued an alert[*28] stating that Japanese emails containing links to malicious URLs that cause Emotet infections had been observed from around December 10, which matches the start of detections in Figure 8. Hence, the traffic detected in Figure 8 is likely accessing URLs in the text of emails designed to spread Emotet.

Next, of the emails thought to be Emotet propagators observed by our SOC, Table 3 summarizes those that contain Japanese text in the subject line. Note that Table 3 only shows the main examples and is not comprehensive.

As Table 3 shows, the subject lines are varied. Some are just a single word, like "Realize", "Help", or "Information", and others purport to be invoices/receipts. Also, around

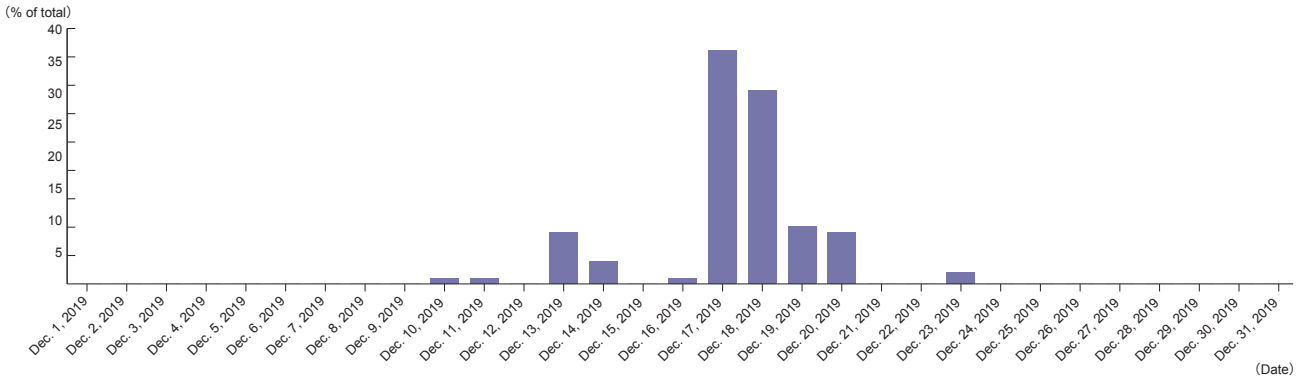

Figure 8: HEUR: Trojan.MSOffice.SAgent Detections as Percentage of Total (Dec. 2019)

Table 3: Suspicious Emails Designed to Spread Emotet with Japanese Text in Subject Line

Subject lines	Attachment filenames	Notes
December bonus	<date>.doc	<date> is the date of receipt in YYYYMMDD format Date, names of people or organizations are appended to the subject line in some cases
[Valid till 23:59 today] Renewal discount coupon issued on amazon.com	<date>_<random alphanumeric string>.doc	
Account credited		
Please issue invoice	<random alphanumeric string> <date>.doc	
Document	<random alphanumeric string>_<date>.doc	
Resending message	<random alphanumeric string>-<date>.doc	
Reminder		
Realize	Bonus payment advice.doc	
Final option	December bonus.doc	
Payment advice	Winter 2019･performance bonus payment.doc	
Help		
Information	Please send invoice <random alphanumeric string>-<date>.doc	
New version	Merry Christmas <date>.doc	
Please attach invoice		
Receipt		

*28 Information-technology Promotion Agency, "Emails designed to propagate a virus called 'Emotet'" (in Japanese, https://www.ipa.go.jp/security/announce/20191202.html#L11).

Black Friday, some subject lines tout discount coupons for online shopping, and attachment filenames contain words to match the season, like “Bonus” or “Christmas”.

■ Countermeasures

As mentioned earlier, Emotet uses information stolen from infected devices to create emails—fake replies etc.—designed to propagate its spread. This may make it difficult for recipients to judge that something is amiss or suspicious based on the sender address or email text. To prevent infections and minimize damage, you should first check your Word settings and disable automatic macro execution if it is on. It is also important not to inadvertently open any attachments or manually enable any macros contained in the attachments that you cannot vet as clean. US-Cert also states that a policy blocking emails with attachments that have filename extensions used by malware or file formats that antivirus software cannot scan is an effective way defend against entry[*29]. It also recommends the use of appropriate permission settings, sender authentication, and the like.

1.4 Conclusion

In this report, we covered prominent security incidents in Japan in 2019 and looked at a number of examples alongside our SOC’s observations. Various security threats beyond these examples are also observed everyday. It is important to properly understand the landscape and address threats, and this effort should not be limited to the incidents and events discussed in Sections 1.2 and 1.3. Some can be addressed with ACL, such the Elasticsearch issues in 1.3.1, while others can be defended against at the individual level by applying vulnerability patches and not casually enabling macros, as discussed in 1.3.3. Our SOC will continue to periodically publish information on security incidents and threats via wizSafe Security Signal (https://wizsafe.iij.ad.jp), and we hope these updates will prove useful in your ongoing security efforts.

Shun Morita
Data Analyst, Security Operations Center, Security Business Department, Advanced Security Division, IIJ

Eisei Honbu
Data Analyst, Security Operations Center, Security Business Department, Advanced Security Division, IIJ

Junya Yamaguchi
Data Analyst, Security Operations Center, Security Business Department, Advanced Security Division, IIJ

*29 CISA, “Increased Emotet Malware Activity” (https://www.us-cert.gov/ncas/current-activity/2020/01/22/increased-emotet-malware-activity).