

Internet Infrastructure Review

Vol.25

November
2014

Infrastructure Security

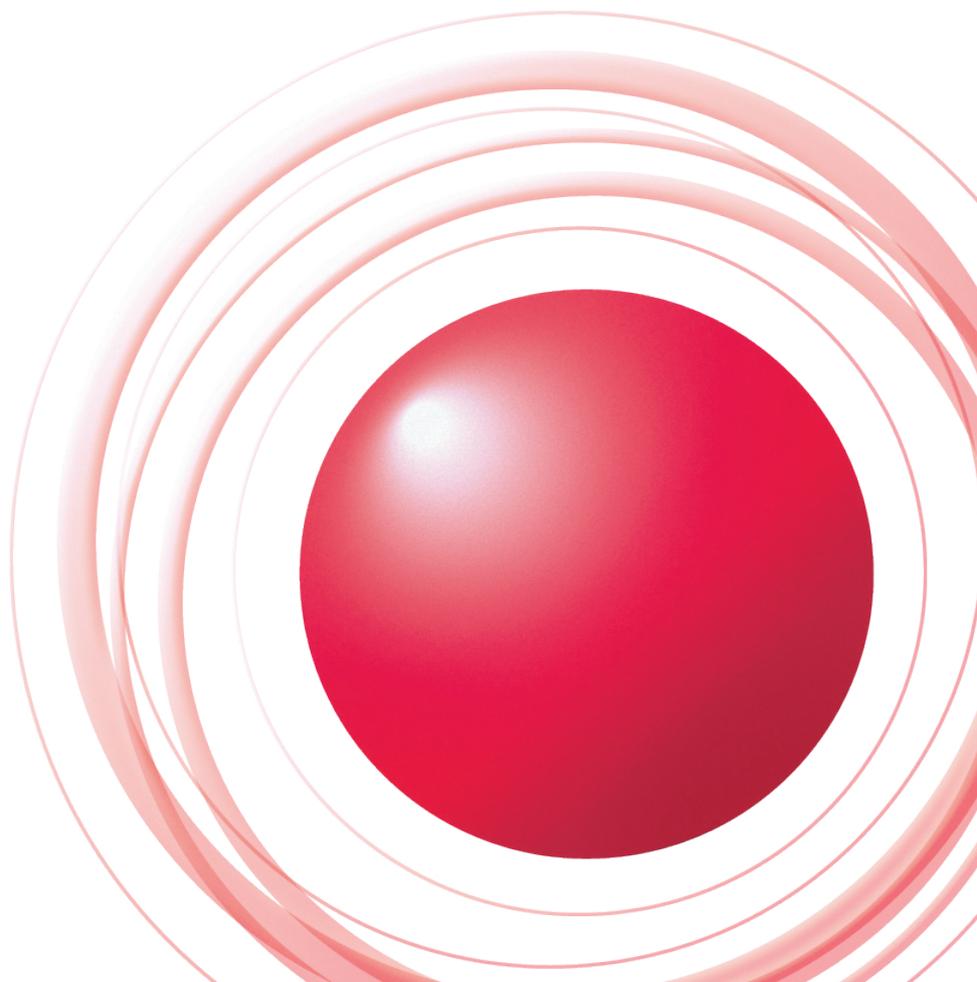
The Shellshock Bash Vulnerability

Content Delivery

The Latest Streaming Technology

Cloud Computing Technology

The Lagopus SDN Software Switch



Executive Summary	3
1. Infrastructure Security	4
1.1 Introduction	4
1.2 Incident Summary	4
1.3 Incident Survey	11
1.3.1 DDoS Attacks	11
1.3.2 Malware Activities	13
1.3.3 SQL Injection Attacks	16
1.3.4 Website Alterations	17
1.4 Focused Research	18
1.4.1 The Shellshock Bash Vulnerability	18
1.4.2 The POODLE Attack	21
1.4.3 The Status of List-Based Attacks and Their Countermeasures	25
1.5 Conclusion	27
2. Content Delivery	28
2.1 Large-Scale Content Delivery	28
2.2 Mobile Topics	29
2.3 The Emergence of H.265 and 4K	29
2.4 4K	30
2.5 MPEG-DASH	30
2.6 Summer Koshien Case Study	32
2.7 Conclusion	33
3. Cloud Computing Technology	34
3.1 SDN and OpenFlow	34
3.2 OpenFlow	35
3.3 Software Switches	36
3.4 The Lagopus SDN Software Switch	36
3.5 Software Switch Bottlenecks	37
3.5.1 Receive Interrupt and Transmission Complete Interrupt Processing	37
3.5.2 Copying Packets to Memory	37
3.5.3 OS Process Scheduling	37
3.5.4 Memory Block Allocation and De-allocation, and Page Faults	37
3.5.5 Exclusive Access to Resources	37
3.6 Resolving Bottlenecks With Intel DPDK	38
3.7 Performance Improvements in Lagopus	38
3.8 Conclusion	39

■ To download current and past issues of the Internet Infrastructure Review in PDF format, please visit the IJ website at <http://www.ij.ad.jp/en/company/development/iir/>.

Executive Summary

The utilization of cloud services by corporate users is increasing year after year. According to the Ministry of Internal Affairs and Communications' FY 2014 Information and Communications White Paper, the ratio of companies using cloud services stood at 33.1% at the end of 2013, which is about a 5% increase over the 28.2% ratio at the end of 2012. This shows that cloud services are on the threshold of going mainstream, with approximately a third of all companies using them in some form or other.

Meanwhile, the white paper also states that 37.4% of companies opting not to use cloud services cite security concerns as their reason. In August, there was an incident in which the private photos of a large number of famous celebrities leaked. These leaks occurred when users' cloud accounts were logged into through account hackings after they had taken photos on a smartphone that were automatically saved to the cloud. Because this was due to accounts being hacked, and did not involve the exploitation of a vulnerability in the system, responsibility for the leaks lay with the victims themselves. According to the white paper, smartphone ownership in Japan has reached 53.5%, and it is conceivable that these will be used for work purposes in more and more cases. To prevent leaks of confidential information such as the example above, it will be increasingly important for each user to have the up-to-date knowledge required to ensure the security of the data they handle, and take appropriate measures to protect it themselves.

This report discusses the results of the various ongoing surveys and analysis activities that IJ carries out to support the Internet infrastructure and enable our customers to continue to use it safely and securely. We also regularly present summaries of technological development as well as important technical information.

In the "Infrastructure Security" section, we give a month-by-month chronological summary of major incidents observed during the three months from July 1 to September 30, 2014, and report on the results of our statistics gathering and analyses for the entire period. We also present our focused research for this period, including discussion of the Bash vulnerability known as Shellshock that was disclosed in September, and a look at a new method for attacking SSLv3 called the POODLE attack that was announced in October. Additionally, we give an overview of the account hacking techniques known as list-based attacks, and examine countermeasures for them.

In the "Content Delivery" section, we explain the current state of video streaming, which has been growing more prevalent in recent years, and is expected to account for around 80% to 90% of consumer traffic worldwide by 2018. We place particular focus on the characteristics of the H.265 codec standard that enables video streaming in ultra-high resolutions such as 4K, as well as MPEG-DASH that makes stable and efficient delivery over the Internet possible. We also present a case study based on the live streaming of the high school baseball summer championship that IJ carried out this year.

In the "Cloud Computing Technology" section, we cover the basics of how SDN and the OpenFlow protocol enable the virtualization and software-based control of networks. We also introduce the Lagopus OpenFlow-compatible software switch that was open sourced this year as part of the O3 strategic R&D project funded by the Ministry of Internal Affairs and Communications to promote SDN. In particular, we discuss the dramatic performance improvements it achieves over conventional software switches, and how it enables packet transfer at a wire-rate of 10GbE.

Through activities such as these, IJ continues to strive towards improving and developing our services on a daily basis while maintaining the stability of the Internet. We will keep providing a variety of solutions that our customers can take full advantage of as infrastructure for their corporate activities.

Author:



Toshiya Asaba

President and CEO, IJ Innovation Institute Inc. President and CEO, Stratosphere Inc. Mr. Asaba joined IJ in its inaugural year of 1992, becoming involved in backbone construction, route control, and interconnectivity with domestic and foreign ISPs. He was named IJ director in 1999, and executive vice president in charge of technical development in 2004. When the IJ Innovation Institute Inc. was founded in June 2008, Mr. Asaba became its president and CEO. When Stratosphere Inc. was founded in April 2012, he also became president and CEO of that organization.

The Shellshock Bash Vulnerability

In this report, we discuss the Shellshock Bash vulnerability, and look at the new POODLE attack method that targets SSLv3. We also examine the occurrence of a series of list-based attacks since last year, and go over some countermeasures for them.

1.1 Introduction

This report summarizes incidents to which IJ responded, based on general information obtained by IJ itself related to the stable operation of the Internet, information from observations of incidents, information acquired through our services, and information obtained from companies and organizations with which IJ has cooperative relationships. This volume covers the period of time from July 1 through September 30, 2014. In this period a number of sites in Japan, including those for media outlets, were affected by DNS hijackings that led to malware infections. A vulnerability was also discovered in GNU Bash, affecting an extensive range of products and services. Unauthorized login incidents due to list-based attacks on online services continued to occur, as did incidents of illegal remittance stemming from the misuse of online banking. There were also SSDP-based DDoS attacks in Japan, coming in the wake of those that used DNS or NTP amplification. Outside Japan, there were sporadic DDoS attacks that exceeded 200 Gbps in bandwidth. These examples show that many security-related incidents continue to occur on the Internet.

1.2 Incident Summary

Here, we discuss the IJ handling and response to incidents that occurred between July 1 and September 30, 2014. Figure 1 shows the distribution of incidents handled during this period*1.

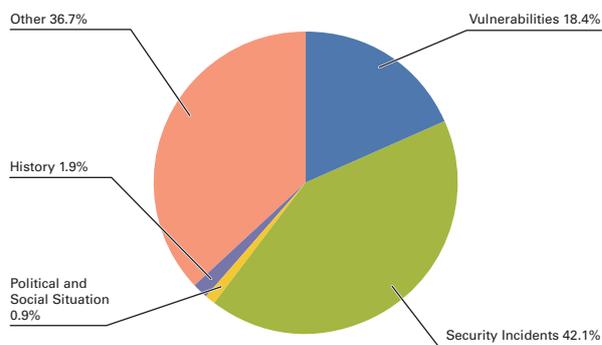


Figure 1: Incident Ratio by Category (July 1 to September 30, 2014)

■ The Activities of Anonymous and Other Hacktivists

Attacks by hacktivists such as Anonymous continued during this period. DDoS attacks and information leaks occurred at government-related and corporate sites in a large number of countries stemming from a variety of incidents and causes. Between July and August, website alterations, DDoS attacks, and information leaks affected a number of government-related sites and private-sector business websites in Israel, in relation to the conflict in Palestinian-controlled Gaza Strip (OpSaveGaza). Attacks were also made between other countries in conflict, such as Russia and Ukraine, India and Indonesia, and India and

*1 Incidents discussed in this report are categorized as vulnerabilities, political and social situation, history, security incidents and other.

Vulnerabilities: Responses to vulnerabilities associated with network equipment, server equipment or software commonly used over the Internet or in user environments.

Political and Social Situations: Responses to incidents related to domestic and foreign circumstances and international events such as international conferences attended by VIPs and attacks originating in international disputes.

History: Historically significant dates; warning/alarms, detection of incidents, measures taken in response, etc., related to attacks in connection with a past historical fact.

Security Incidents: Unexpected incidents and related responses such as wide propagation of network worms and other malware; DDoS attacks against certain websites.

Other: Security-related information, and incidents not directly associated with security problems, including highly concentrated traffic associated with a notable event.

Pakistan. Between August and September, there were protests against the Pakistani government in the form of website alterations, DDoS attacks, and information leaks through SQL injections that affected a number of government-related sites (OpPakistan). In September, there were DDoS attacks and leaks of email account information that affected websites of the Chinese government and the Hong Kong executive branch, in relation to demonstrations about issues with the electoral system in Hong Kong (OpHongKong).

Other attacks by hacktivists such as Anonymous continued on government and government-related websites around the world. Unknown attackers claiming affiliation with the Syrian Electronic Army continued to hijack SNS accounts and deface websites, with affected accounts including those for the Israel Defense Forces.

■ Vulnerabilities and their Handling

During this period, fixes were released for Microsoft's Windows^{*2*3} and Internet Explorer^{*4*5*6}. Updates were also made to Adobe Systems' Adobe Flash Player, Adobe Reader, and Acrobat. A quarterly update was provided for Oracle's Java SE, fixing many vulnerabilities. Several of these vulnerabilities were exploited before patches were released. Regarding server applications, a quarterly update was released for a number of Oracle products, including the Oracle database server, fixing many vulnerabilities.

A number of vulnerabilities, including those that allowed elevation of privileges, were discovered and fixed in Cisco Unified Communications Domain Manager. One of the vulnerabilities fixed involved the use of the default SSH private key. Using the default SSH private key on a device without updating it leads to a high risk of exploitation by third parties with the same private key, so this should be avoided^{*7}.

A vulnerability (CVE-2014-6271) that could allow arbitrary OS commands to be executed was discovered and fixed in the Bash shell used in Unix-based OSes such as Linux. This vulnerability, which was given the name Shellshock, became a problem because it affected a large number of applications and devices, such as Web servers that run CGI scripts, routers, and gateway products. See "1.4.1 The Shellshock Bash Vulnerability" for more information about this issue.

■ Unauthorized Login Through Identity Fraud

Since last year there have been many attempts to steal user IDs and passwords, and log in without authorization presumably using lists of these IDs and passwords. These attempts continued in the current survey period. A variety of sites were targeted in these attacks, including survey sites, e-commerce sites, support sites for logistics companies, sites for mobile phone companies, and SNS. In a number of these incidents tangible damage was caused, such as the exchange of site points for gift points on other sites without authorization. As this demonstrates, the threat from list-based attacks still persists. A range of companies are taking steps to enhance the security of authentication functions for Web services and applications, such as adding two-step authentication or relaxing restrictions on password length and usable characters. Because problems have also been identified with users reusing the same password, or setting simple passwords that are easy to guess^{*8}, users must also be sure to follow safe practices. See "1.4.3 The Status of List-Based Attacks and Their Countermeasures" for more information.

*2 "Microsoft Security Bulletin MS14-038 - Critical: Vulnerability in Windows Journal Could Allow Remote Code Execution (2975689)" (<https://technet.microsoft.com/library/security/ms14-038>).

*3 "Microsoft Security Bulletin MS14-043 - Critical: Vulnerability in Windows Media Center Could Allow Remote Code Execution (2978742)" (<https://technet.microsoft.com/library/security/ms14-043>).

*4 "Microsoft Security Bulletin MS14-037 - Critical: Cumulative Security Update for Internet Explorer (2975687)" (<https://technet.microsoft.com/library/security/ms14-037>).

*5 "Microsoft Security Bulletin MS14-051 - Critical: Cumulative Security Update for Internet Explorer (2976627)" (<https://technet.microsoft.com/library/security/ms14-051>).

*6 "Microsoft Security Bulletin MS14-052 - Critical: Cumulative Security Update for Internet Explorer (2977629)" (<https://technet.microsoft.com/library/security/ms14-052>).

*7 See the IJ Security Diary post, "The Issue of Many Public Keys Unintentionally Sharing Private Keys with Other Sites" (<https://sect.ij.ad.jp/d/2012/08/109998.html>) (in Japanese) for information on the risks of using the same private key.

*8 For example, see the IPA "Report on the 'Survey of Online Personal Authentication Systems'" (<http://www.ipa.go.jp/security/fy26/reports/ninsho/index.html>) (in Japanese).

July Incidents

1	V 3rd: A number of vulnerabilities, including those that allowed elevation of privileges, were discovered and fixed in Cisco Unified Communications Domain Manager. "Multiple Vulnerabilities in Cisco Unified Communications Domain Manager" (http://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20140702-cucdm).
2	
3	V 9th: Microsoft published their Security Bulletin Summary for July 2014, and released six updates, including the two critical updates MS14-037 and MS14-038, as well as three important updates. "Microsoft Security Bulletin Summary for July 2014" (https://technet.microsoft.com/library/security/ms14-jul).
4	
5	V 9th: A number of vulnerabilities in Adobe Flash Player that could allow arbitrary code execution were discovered and fixed. "APSB14-17: Security updates available for Adobe Flash Player" (http://helpx.adobe.com/security/products/flash-player/apsb14-17.html).
6	S 9th: A correspondence education company announced that information on 20,700,000 of their customers had leaked to third parties including list traders. Later, in a final report published in September, it was revealed that the personal information of 48,580,000 customers had leaked.
7	S 9th: It came to light that certificates for a number of Google and Yahoo! domains had been issued without authorization at the National Informatics Center of India (NIC), which operates an intermediate certificate authority affiliated with the root certificate authority of the government of India. As a result, these certificates were revoked in a number of browsers. It is thought that this happened when the certificate issuing process was compromised. Google Online Security Blog, "Maintaining digital certificate security" (http://googleonlinesecurity.blogspot.jp/2014/07/maintaining-digital-certificate-security.html).
8	
9	
10	O 15th: The Amended Act Prohibiting Child Prostitution and Pornography that was passed in June came into effect. However, regarding the provisions of Article 7-1 of the amended act, which details punishment for the possession of child pornography for the purpose of fulfilling personal sexual curiosity, a grace period of one year from enactment in which such material could be disposed of appropriately was given in an additional clause. This article will now apply from July 15, 2015. See the following Ministry of Justice explanation for more details. "The bill for amending part of the Act on Punishment of Activities Relating to Child Prostitution and Child Pornography, and the Protection of Children" (http://www.moj.go.jp/keiji1/keiji11_00008.html) (in Japanese).
11	
12	
13	O 15th: The Ministry of Internal Affairs and Communications published its "Information and Communications in Japan 2014" white paper, which examines the current status of information and communications in Japan, as well as trends in information and communications policy. Information and Communications White Paper Site (http://www.soumu.go.jp/johotsusintokei/whitepaper/index.html) (in Japanese).
14	
15	V 16th: Oracle released their quarterly scheduled update for a number of products including Oracle, fixing a total of 113 vulnerabilities, including 20 in Java SE. "Oracle Critical Patch Update Advisory - July 2014" (http://www.oracle.com/technetwork/topics/security/cpujul2014-1972956.html).
16	O 16th: The IPA officially established its Cyber Rescue Squad to support organizations affected by targeted attacks. Their role will be to limit damages, deter and reduce reoccurrence, and implement swift countermeasures to prevent a chain reaction effect. "Press Release - Cyber Rescue Squad Established to Support Countermeasures to Targeted Attacks" (https://www.ipa.go.jp/about/press/20140716_1.html) (in Japanese).
17	
18	S 17th: Regarding an incident in which customer information was leaked from a correspondence education company, a former temporary employee of a subcontractor was arrested on suspicion of violating the Unfair Competition Prevention Act (copying of trade secrets).
19	
20	S 18th: The Telecom Information Sharing and Analysis Center Japan announced it would issue alerts to users affected by a type of malware that targets Internet banking (Gameover Zeus) through its public-private project to support malware countermeasures in Japan (ACTIVE), as part of an international operation to remove infections. "Regarding Alerts for Users Infected with Malware That Targets Internet Banking" (https://www.telecom-isac.jp/news/news20140718.html) (in Japanese).
21	
22	O 22nd: The Council for Stable Operation of the Internet published the third edition of its "Guidelines for Dealing with High Volume Communications and Privacy at Telecommunications Carriers." "Guidelines for Dealing with High Volume Communications and Privacy at Telecommunications Carriers (Third Edition)" (http://www.jaipa.or.jp/other/mtcs/guideline_v3.pdf) (in Japanese).
23	
24	O 23rd: The Ministry of Internal Affairs and Communications announced it would hold a "Workshop on Personal Data Held by Government Institutions, etc." to investigate and evaluate the possible use and protection of personal data in the possession of government institutions and independent administrative agencies. "Workshop on Personal Data Held by Government Institutions, etc." (http://www.soumu.go.jp/main_sosiki/kenkyu/gyousei_personal/index.html) (in Japanese).
25	
26	
27	S 24th: The European Central Bank announced its website had been compromised, and that personal data including the email addresses of event registrants had leaked. This incident came to light when an anonymous email demanding money in return for the leaked data was received. See the following European Central Bank announcement for more details. "24 July 2014 - ECB announces theft of contact information" (https://www.ecb.europa.eu/press/pr/date/2014/html/pr140724.en.html).
28	
29	S 25th: A number of companies and government institutions in Japan issued a series of alerts regarding websites that imitated their own sites.
30	V 29th: Fixes were made to a vulnerability in a number of IP cameras that allowed authentication to be bypassed, leading to the possibility of arbitrary operations being executed by obtaining device configuration details such as authentication information. JVN, "JVND-2014-000087 Multiple I-O DATA IP Cameras vulnerable to authentication bypass" (http://jvndb.jvn.jp/en/contents/2014/JVND-2014-000087.html).
31	

[Legend]

**Vulnerabilities****Security Incidents****Political and Social Situation****History****Other**

*Dates are in Japan Standard Time

■ Website Alterations and Redirection to Malware

During this survey period, there continued to be many incidents in which websites were altered to redirect visitors to malicious software. There were website alterations and resulting redirection to malware at a tour company, an automobile sales company, and a security firm. In addition to redirection to malware sites, in an incident in which the website of an independent administrative agency was altered, there were attempts to exploit it as a phishing site*⁹. A text editing software support site was also altered, redirecting users accessing it to another site where malware was installed. After the incident came to light, it was revealed that the content of a legitimate update file was also altered in an attempt to get users to install a virus-infected file. We believe this kind of malware infection activity that exploits update systems for legitimate software will continue in the future.

■ Attacks Based on Political and Social Situation and Historical Context

During this period each year there are incidents related to historical dates in the Pacific War, as well as Takeshima and the Senkaku Islands. We stayed vigilant, as this year it was expected that the websites of a number of government agencies and private-sector businesses in Japan would be subject to alterations through compromise via SQL injections and brute force attacks, as well as DDoS attacks, in relation to these sensitive issues. However, although there was an increase in unauthorized access through SQL injection attacks in some cases*¹⁰, and website alterations also took place*¹¹, no large-scale attacks were confirmed. IIJ's observations showed slightly more DDoS attacks than usual, but the number and scale of attacks was significantly lower than the same period in previous years.

■ DDoS Attacks

A number of large-scale DDoS attacks occurred during this period. In August, an individual or group calling themselves the Lizard Squad attacked a number of game-related servers, including those for PSN, Xbox Live, and League of Legends. The incident affecting PSN is thought to have involved an NTP reflection attack of 263.35 Gbps. In addition to DDoS attacks, the perpetrator also made claims online that they had placed a bomb on the plane an executive was flying in, leading to the flight being diverted. Attacks on game-related servers thought to be the work of this attacker also continued intermittently in September. In Japan, periodic DDoS attacks were made on the DNS servers of a number of ISPs in late May, and these continued through to July at some ISPs. In September, a high school student was referred to prosecutors on the charge of obstruction of business by damaging a computer, having been suspected of attacking game servers in March using an overseas DDoS attack service.

■ Malware Infections and Information Leaks at Companies

During this survey period, there were ongoing incidents of large-scale leaks of data including customer information due to corporate systems being infected with malware. In particular, it was announced in August that 51 domestic branches of a major U.S. logistics company had been infected with malware, leading to the possibility that details such as customer credit card information had been leaked. There were also incidents in Japan, including malware infections of terminals at an airline company, which could have resulted in customer information being sent to an external party and leaked. In September, there was an incident at a major home improvement center company in which the details of approximately 56 million payment cards and email addresses may have leaked. It is thought that a variant of the malware targeted at POS systems that has been active since last year was used in this incident, and data including credit card information could have leaked. Information leaks due to malware also took place at a number of other companies, including hospitals and retailers. Due to POS malware in particular causing multiple large-scale information leaks since around the end of last year, US-CERT issued a number of alerts*¹². New POS malware variants continue to appear*¹³, so ongoing vigilance will be necessary.

*⁹ For example, see the following National Research Institute for Earth Science and Disaster Prevention announcement. "Regarding Alterations to the NIED Website" (http://www.bosai.go.jp/press/2014/pdf/20140811_01.pdf) (in Japanese).

*¹⁰ For example, it has been reported in IBM's Tokyo SOC Report that attacks such as SQL injections from specific countries are on the rise. "Attack trends around September 18 when the Liutiaohu Incident took place" (https://www-304.ibm.com/connections/blogs/tokyo-soc/entry/201409attackfromcn?lang=en_us) (in Japanese).

*¹¹ See the following F-Secure blog post for more information. "Sept. 18 cyber attacks (follow-up)" (<http://blog.f-secure.jp/archives/50734688.html>) (in Japanese).

*¹² US-CERT issues ongoing alerts, such as "TA14-002A: Malware Targeting Point of Sale Systems" (<https://www.us-cert.gov/ncas/alerts/TA14-002A>) in January and "Alert (TA14-212A) Backoff Point-of-Sale Malware" (<https://www.us-cert.gov/ncas/alerts/TA14-212A>) in August.

*¹³ For example, see the Trend Micro blog post "2014 - An Explosion of Data Breaches and PoS RAM Scrapers" (<http://blog.trendmicro.com/trendlabs-security-intelligence/2014-an-explosion-of-data-breaches-and-pos-ram-scrapers/>) for more information.

August Incidents

1	S 1st: US-CERT issued an alert regarding a newly discovered type of malware called Backoff that targets POS systems. "Alert (TA14-212A) Backoff Point-of-Sale Malware" (https://www.us-cert.gov/ncas/alerts/TA14-212A).
2	O 1st: Microsoft released the Enhanced Mitigation Experience Toolkit (EMET) 5.0, which is a security tool for mitigating application vulnerabilities. See the following TechNet Blogs post for more information. "Announcing EMET 5.0" (http://blogs.technet.com/b/srd/archive/2014/07/31/announcing-emet-v5.aspx).
3	
4	S 2nd: The Mozilla Developer Network announced that 76,000 MDN user email addresses and 4,000 encrypted user passwords may have leaked due to a database dump file being mistakenly made available to the public.
5	Mozilla Developer Network, "MDN Database Disclosure" (https://blog.mozilla.org/security/2014/08/01/mdn-database-disclosure/).
6	
7	S 4th: There was an incident of unauthorized access at Gamma International, the provider of the FinSpy (FinFisher) commercial surveillance software allegedly used by government institutions around the world for intelligence gathering activities. As a result, a total of 40 GB of internal documents and source code was released.
8	
9	S 8th: A U.S. security company reported that BGP hijacking attacks had been used to redirect traffic destined for a virtual currency mining pool to a fake mining pool between February and May 2014. It is thought that 19 ISPs were affected by these attacks, and the attacker may have profited by as much as \$83,000.
10	See the following Dell SecureWorks blog post for more information. "BGP Hijacking for Cryptocurrency Profit" (http://www.secureworks.com/cyber-threat-intelligence/threats/bgp-hijacking-for-cryptocurrency-profit/).
11	
12	S 12th: In the United States, a phenomenon that caused Internet speeds to drop and communications to become unstable occurred. This is thought to have been due to the advertised BGP routing information exceeding the maximum BGP routing table size of 512 k on older routers.
13	See the following BGPmon.net blog post for more information. "What caused today's Internet hiccup" (http://www.bgpmn.net/what-caused-todays-internet-hiccup/).
14	
15	V 13th: Microsoft published their Security Bulletin Summary for August 2014, and released two critical updates, MS14-043 and MS14-051, as well as seven important updates.
16	"Microsoft Security Bulletin Summary for August 2014" (https://technet.microsoft.com/library/security/ms14-aug).
17	V 13th: A number of vulnerabilities in Adobe Flash Player that could allow arbitrary code execution were discovered and fixed.
18	"APSB14-18: Security updates available for Adobe Flash Player" (http://helpx.adobe.com/security/products/flash-player/apsb14-18.html).
19	V 13th: A vulnerability in Adobe Reader and Acrobat that could allow remote arbitrary code execution was discovered and fixed.
20	"Security updates available for Adobe Reader and Acrobat" (http://helpx.adobe.com/security/products/reader/apsb14-19.html).
21	S 13th: It was announced that the Japanese and Simplified Chinese support sites for a text editor application had been altered, and traces of attempts to steal user names, passwords, and IP addresses were found. The website was subsequently altered again on August 18, and the update checking function of the shareware product was used to install malicious files in some cases.
22	
23	V 15th: Microsoft published a recommendation to uninstall the MS14-045 update released in August 2014 as a preventative measure, even if no issues were experienced, due to the possibility of it causing abnormal shutdowns or boot failures when applied.
24	For more information, see the following explanation from the Japan security team's official blog "[Issue found after release] Installing the update released on August 13, 2014 may cause issues" (http://blogs.technet.com/b/jpsecurity/archive/2014/08/16/2982791-knownissue3.aspx) (in Japanese).
25	S 21st: U.S. UPS announced it had confirmed malware infections at 51 of its domestic branches, and that data such as the credit card information of its customers may have leaked.
26	See the following United Parcel Service of America, Inc. announcement for details. "The UPS Store, Inc. Notifies Customers Of Potential Data Compromise and Incident Resolution" (http://www.pressroom.ups.com/Press+Releases/Archive/2014/Q3/The+UPS+Store%2C+Inc.+Notifies+Customers+Of+Potential+Data+Compromise+and+Incident+Resolution).
27	
28	S 25th: The PlayStation Network (PSN) and Sony Entertainment Network (SEN) were targeted in DDoS attacks carried out by an unknown party, causing major faults.
29	See the following PlayStation.Blog post for more information. "Update: PlayStation Network is Back Online" (http://blog.us.playstation.com/2014/08/24/playstation-network-update-2/).
30	
31	V 28th: Microsoft fixed and rereleased MS14-045, which had caused issues such as abnormal shutdowns and boot failures after being applied.
	"Microsoft Security Bulletin MS14-045 - Important: Vulnerabilities in Kernel-Mode Drivers Could Allow Elevation of Privilege (2984615)" (https://technet.microsoft.com/library/security/ms14-045).

[Legend]



Vulnerabilities



Security Incidents



Political and Social Situation



History



Other

*Dates are in Japan Standard Time

■ Government Agency Initiatives

Government agency initiatives included holding the 40th assembly of the government's Information Security Policy Council. The council published the "Annual Report on Cyber Security Policy (Fiscal 2013)," their first annual report that summarizes the overall status of cyber security in Japan for fiscal 2013, including the status of relevant policy at government institutions, critical infrastructure providers, and ministries, as well as related documents. Decisions were also made on "Cyber Security 2014," which compiles information such as the cyber security policies at each ministry for the current fiscal year*14.

Based on the "Policy Outline of the Institutional Revision for Utilization of Personal Data" that was published in June of this year, the Ministry of Internal Affairs and Communications held a "Workshop on Personal Data Held by Government Institutions, etc.," to investigate and evaluate the possible use and protection of personal data in the possession of government institutions and independent administrative agencies. They also published their "Location Information Privacy Report - Appropriately Protecting Location Information Privacy While Allowing Its Social Use and Application," which summarizes the proper handling of the acquisition and use of location information, as well as its provision to third parties*15. The aim of this report is to appropriately protect the confidentiality of communications, personal information, and privacy regarding the location information handled by telecommunications carriers, while promoting its social use and application, including for business purposes.

The National Police Agency published their White Paper on Police for 2014, which stated they had made a record number of arrests for cybercrimes, and indicated a big surge in the number of illegal remittances via Internet banking. The report also discussed initiatives such as their efforts to promote public and private coordination with regard to threats in cyberspace.

■ Other

Since June, countermeasures have been carried out for the Gameover Zeus malware that steals online banking information*16. In July, it was announced that information regarding terminals infected with this malware obtained through the operation would be used to provide ISPs with the details of infected parties, and alerts would be issued using this information*17. Additionally, progress is being made towards building a framework for international collaboration. One example of this is the establishment of a Joint Cybercrime Action Taskforce (J-CAT) by the European Police Office (Europol) in September*18. The task force covers a number of regions including Europe, the United States, and Canada, and its purpose is to deal with international cybercrime such as the examples above.

In July, a range of organizations and companies issued alerts regarding counterfeit websites*19. In August, government institutions also issued alerts regarding similar incidents. It is thought that these incidents originated from a Web proxy service, and did not target specific companies or organizations. Similar examples include those originating from a website conversion proxy provided by a mobile telecommunications carrier in China last year*20. Many other similar services are available, but because it is often not clear who the provider is, caution must be exercised when using them*21.

Also in July, the leak of customer information from a correspondence education company came to light. The incident was uncovered when another company sent direct mail based on the list of names obtained. A former temporary employee of a subcontractor of the affected company was arrested and charged with violating the Unfair Competition Prevention Act (copying of trade secrets), on suspicion of taking the data without authorization and selling it to list traders, etc.

*14 National Information Security Center, "Information Security Policy Council - 40th Assembly" (July 10, 2014) (<http://www.nisc.go.jp/conference/seisaku/index.html#seisaku40>) (in Japanese).

*15 Ministry of Internal Affairs and Communications, "Investigative Commission on the Handling of Location Information in Emergencies' 'Location Information Privacy Report - Appropriately Protecting Location Information Privacy While Allowing Its Social Use and Application' Report Published" (http://www.soumu.go.jp/menu_news/s-news/01kiban08_02000144.html) (in Japanese).

*16 Department of Justice, "U.S. Leads Multi-National Action Against 'Gameover Zeus' Botnet and 'Cryptolocker' Ransomware, Charges Botnet Administrator" (<http://www.justice.gov/opa/pr/2014/June/14-crm-584.html>).

*17 Telecom Information Sharing and Analysis Center Japan, "Regarding Alerts for Users Infected with Malware That Targets Internet Banking" (<https://www.telecom-isac.jp/news/news20140718.html>) (in Japanese).

*18 See the following European Central Bank announcement for more details. "EXPERT INTERNATIONAL CYBERCRIME TASKFORCE IS LAUNCHED TO TACKLE ONLINE CRIME" (<https://www.europol.europa.eu/content/expert-international-cybercrime-taskforce-launched-tackle-online-crime>).

*19 For example, see the following Osaka Prefectural Police announcement, "Urgent: Beware of Imitations of the Osaka Prefectural Police Website!" (http://www.police.pref.osaka.jp/15topics/caution_domain.html) (in Japanese).

*20 See the following JPCERT Coordination Center tweet for more information (<https://twitter.com/jpcert/status/322282948554530816>) (in Japanese).

*21 See the following Trend Micro blog post for more information about this incident, "What lessons can be learned from the imitative site confusion caused by proxy avoidance systems?" (<http://blog.trendmicro.co.jp/archives/9713>) (in Japanese).

September Incidents

1	S 1st: In the U.S., the personal photos of celebrities including Hollywood actresses were posted to a message board. It is thought that this incident originated from unauthorized access to the iCloud accounts of a number of celebrities that had leaked. Apple published the following report on their investigation. "Apple Media Advisory Update to Celebrity Photo Investigation" (http://www.apple.com/pr/library/2014/09/02Apple-Media-Advisory.html).
2	
3	O 1st: The Ministry of Internal Affairs and Communications published its "Telecommunications Service Accident Status Report (2013)," which summarized the status of telecommunications accident reports for 2013. "Telecommunications Service Accident Status Report (2013)" (http://www.soumu.go.jp/menu_news/s-news/01kiban05_02000072.html) (in Japanese).
4	
5	S 3rd: A large-scale leak of card information occurred at major U.S. retailer The Home Depot. These indicated that the total amount of money involved had surpassed the second half of 2013, and the scope of these incidents had spread to regional banks, credit unions, and credit associations. They also showed there was a sharp rise in incidents affecting corporate accounts. "Status of Incidents of Illegal Remittance Related to Internet Banking in the First Half of 2014" (http://www.npa.go.jp/cyber/pdf/H260904_banking.pdf) (in Japanese).
6	
7	O 4th: The National Police Agency announced its statistics on unauthorized remittance crimes via Internet banking for the first half of 2014. These indicated that the total amount of money involved had surpassed the second half of 2013, and the scope of these incidents had spread to regional banks, credit unions, and credit associations. They also showed there was a sharp rise in incidents affecting corporate accounts. "Status of Incidents of Illegal Remittance Related to Internet Banking in the First Half of 2014" (http://www.npa.go.jp/cyber/pdf/H260904_banking.pdf) (in Japanese).
8	
9	
10	V 10th: Microsoft published their Security Bulletin Summary for September 2014, and released the MS14-052 critical update, as well as three important updates. "Microsoft Security Bulletin Summary for September 2014" (https://technet.microsoft.com/library/security/ms14-sep).
11	
12	V 10th: A number of vulnerabilities in Adobe Flash Player that could allow unauthorized termination and arbitrary code execution were discovered and fixed. "APSB14-21: Security updates available for Adobe Flash Player" (http://helpx.adobe.com/security/products/flash-player/apsb14-21.html).
13	
14	O 11th: The National Police Agency published its "Report on Cyberspace Threats for the First Half of 2014," which gave an overview of trends in cybercrime for the first half of 2014. "Report on Cyberspace Threats for the First Half of 2014" (http://www.npa.go.jp/kanbou/cybersecurity/H26_kami_jousei.pdf) (in Japanese).
15	
16	V 17th: A number of vulnerabilities in Adobe Reader and Acrobat that could allow unauthorized termination and remote arbitrary code execution were discovered and fixed. "APSB14-20: Security updates available for Adobe Reader and Acrobat" (http://helpx.adobe.com/security/products/reader/apsb14-20.html).
17	
18	S 18th: A number of DDoS attacks were carried out on game servers, and a high school student was referred to prosecutors on suspicion of obstructing business by damaging a computer to interfere with game company operations.
19	
20	P 18th: Attacks often occur around this day each year for historical reasons. However, although there were small-scale attacks this year, no organized attacks were observed.
21	
22	O 22nd: LINE Corporation made it compulsory to set a PIN number (four digits) for smartphone versions of the LINE app, as a countermeasure to limit damages from frequent LINE account hijackings. See the following official LINE blog post for more information "[Important] PIN numbers will be made compulsory to prevent the spread of unauthorized login (hijacking) incidents" (http://official-blog.line.me/ja/archives/1009539887.html) (in Japanese).
23	
24	S 24th: It was announced that unauthorized access had taken place at an airline company due to a malware infection, leading to the possibility that the personal information of up to 730,000 members had leaked.
25	
26	V 25th: A vulnerability in Bash that allowed arbitrary code execution was discovered and fixed. This fix was subsequently found to be insufficient, so further fixes were made for a number of vulnerabilities. CERT/CC, "Vulnerability Note VU#252743 - GNU Bash shell executes commands in exported functions in environment variables" (http://www.kb.cert.org/vuls/id/252743).
27	
28	V 26th: A major cloud provider made news when they carried out large-scale maintenance due to an undisclosed vulnerability in Xen. On October 2, it was announced that this was to deal with CVE-2014-7188. See the following Xen Project Blog post for more information. "XSA-108: Additional Information from the Xen Project" (https://blog.xenproject.org/2014/10/02/xsa-108-additional-information-from-the-xen-project-2/).
29	
30	S 26th: It was announced that unauthorized login thought to be a password list attack had targeted the support site for a logistics company, and the personal information of some members may have leaked through being viewed there. A similar incident took place at another logistics company on the 28th.

[Legend]

V Vulnerabilities**S** Security Incidents**P** Political and Social Situation**H** History**O** Other

*Dates are in Japan Standard Time

Revisions were made to the “Guidelines for Dealing with High Volume Communications and Privacy at Telecommunications Carriers,” which are designed to help telecommunications carriers identify and deal with high volume communications such as DoS. The guidelines were then published by five telecommunications carrier organizations*²². The revisions were made based on the initial report of the Ministry of Internal Affairs and Communications’ Workshop on the Appropriate Way to Handle Cyber Attacks in the Telecommunications Business, which was published in April*²³. They incorporated measures for dealing with new threats such as DNS amplification attacks.

Also, from early September through to October, the .com domains used by a number of organizations in Japan were targeted in DNS hijackings, leading to attempts to install malicious software*²⁴. Because these incidents involved the name server information registered to the registry being rewritten by some means to redirect traffic to a fraudulent website prepared by the attacker, JPCERT/CC and JPRS issued alerts*²⁵.

1.3 Incident Survey

1.3.1 DDoS Attacks

Today, DDoS attacks on corporate servers are almost a daily occurrence, and the methods involved vary widely. However, most of these attacks are not the type that utilizes advanced knowledge such as that of vulnerabilities, but rather cause large volumes of unnecessary traffic to overwhelm network bandwidth or server processes for the purpose of hindering services.

■ Direct Observations

Figure 2 shows the circumstances of DDoS attacks handled by the IJ DDOS Protection Service between July 1 and September 30, 2014.

This information shows traffic anomalies judged to be attacks based on IJ DDOS Protection Service standards. IJ also responds to other DDoS attacks, but these incidents are excluded from the figure due to the difficulty in accurately ascertaining the facts of each situation.

There are many methods that can be used to carry out a DDoS attack, and the capacity of the environment attacked (bandwidth and server performance) will largely determine the degree of impact. Figure 2 categorizes DDoS attacks into three types:

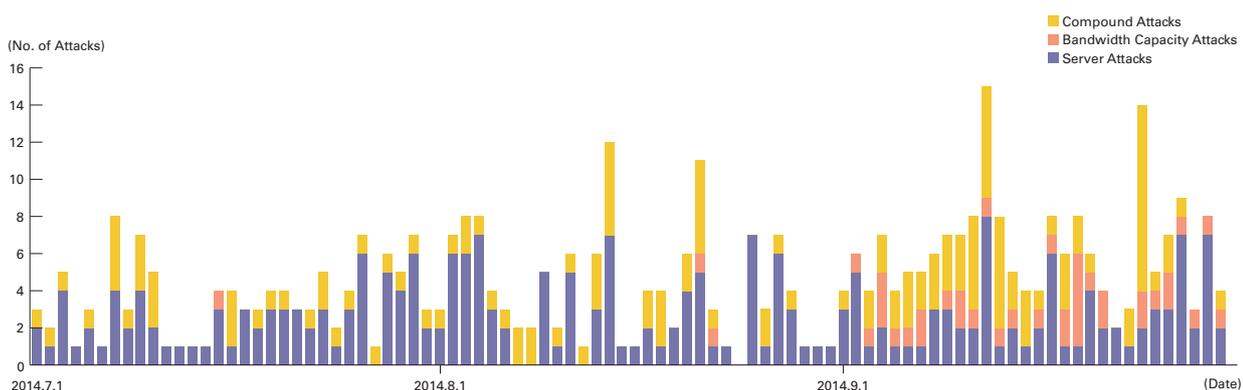


Figure 2: Trends in DDoS Attacks

*22 The Council for Stable Operation of the Internet is comprised of five industry associations with links to the telecommunications business. Refer to the website below for details on the establishment of these guidelines. Japan Internet Provider’s Association (JAIPA) “Revision to Guidelines for Dealing with High Volume Communications and Privacy at Telecommunications Carriers” (<http://www.jaipa.or.jp/topics/?p=695>) (in Japanese).

*23 Ministry of Internal Affairs and Communications, “Initial Report of the Workshop on the Appropriate Way to Handle Cyber Attacks in the Telecommunications Business’ and Results of Request for Public Comment Published” (http://www.soumu.go.jp/menu_news/s-news/01ryutsu03_02000074.html) (in Japanese).

*24 Volexity Blog, “Democracy in Hong Kong Under Attack” (<http://www.volexity.com/blog/?p=33>).

*25 Japan Registry Services (JPRS), “(Urgent) Regarding domain name hijackings caused by unauthorized rewriting of registered information and their countermeasures (published November 5, 2014) (<http://jprs.jp/tech/security/2014-11-05-unauthorized-update-of-registration-information.html>). (in Japanese).

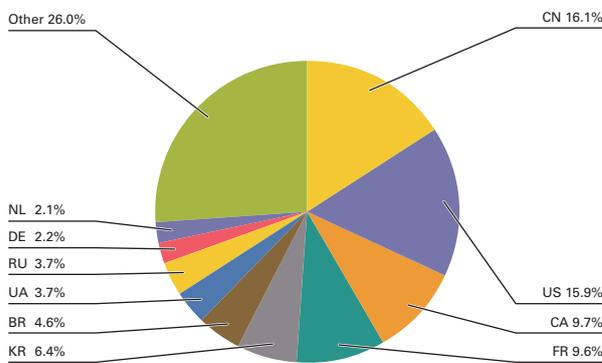
attacks on bandwidth capacity^{*26}, attacks on servers^{*27}, and compound attacks (several types of attacks on a single target conducted at the same time).

During the three months under study, IIJ dealt with 340 DDoS attacks. This averages to 3.7 attacks per day, indicating a decrease in the average daily number of attacks compared to our prior report. Server attacks accounted for 71.5% of all incidents, while compound attacks accounted for 16.8%, and bandwidth capacity attacks 11.8%.

The largest attack observed during the period under study was classified as a compound attack, and resulted in 4.88 Gbps of bandwidth using up to 486,000 pps packets.

Of all attacks, 90.9% ended within 30 minutes of commencement, 9.1% lasted between 30 minutes and 24 hours, and none lasted over 24 hours. The longest sustained attack was a compound attack that lasted for 17 hours and 36 minutes.

Each year during this period, many DDoS attacks are observed around historic dates. From the beginning of September there was an increase in DDoS attacks, and a change was also seen in attack trends, but because these were not organized attacks, we couldn't link them to any particular cause.



In most cases, we observed an extremely large number of IP addresses, whether domestic or foreign. We believe this is accounted for by the use of IP spoofing^{*28} and botnet^{*29} usage as the method for conducting DDoS attacks.

Figure 3: DDoS Attack Targets by Country According to Backscatter Observations

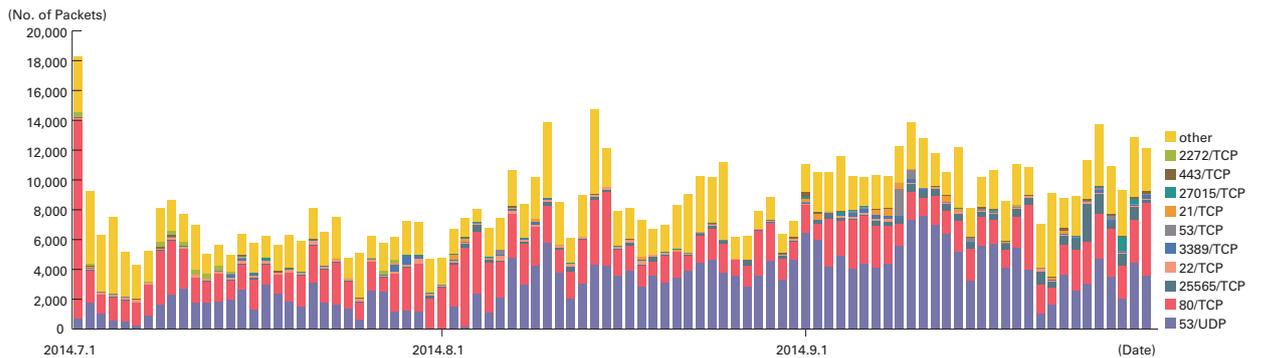


Figure 4: Observations of Backscatter Caused by DDoS Attacks (Observed Packets, Trends by Port)

^{*26} Attack that overwhelms the network bandwidth capacity of a target by sending massive volumes of larger-than-necessary IP packets and fragments. The use of UDP packets is called a UDP flood, while the use of ICMP packets is called an ICMP flood.

^{*27} TCP SYN flood, TCP connection flood, and HTTP GET flood attacks. TCP SYN flood attacks send mass volumes of SYN packets that signal the start of TCP connections, forcing the target to prepare for major incoming connections, causing the wastage of processing capacity and memory. TCP connection flood attacks establish mass volumes of actual TCP connections. HTTP GET flood attacks establish TCP connections on a Web server, and then send mass volumes of HTTP GET protocol commands, wasting processing capacity and memory.

^{*28} Misrepresentation of a sender's IP address. Creates and sends an attack packet that has been given an address other than the actual IP address of the attacker to make it appear as if the attack is coming from a different location, or from a large number of individuals.

^{*29} A "bot" is a type of malware that institutes an attack after receiving a command from an external C&C server. A network constructed of a large number of bots acting in concert is called a botnet.

■ Backscatter Observations

Next we present our observations of DDoS attack backscatter using the honeypots^{*30} set up by the MITF, a malware activity observation project operated by IIJ^{*31}. By monitoring backscatter it is possible to detect some of the DDoS attacks occurring on external networks as a third party without any interposition.

For the backscatter observed between July 1 and September 30, 2014, Figure 3 shows the sender's IP addresses classified by country, and Figure 4 shows trends in packet numbers by port.

The port most commonly targeted by the DDoS attacks observed was the 53/UDP port used for DNS, accounting for 36.2% of the total during the target period. This was followed by 80/TCP used for Web services at 27.3%, so the top two ports accounted for 63.5% of the total. Attacks were also observed on 53/TCP used for DNS, 22/TCP used for SSH, 3389/TCP used for remote desktop, 21/TCP used for FTP, and 443/TCP used for HTTPS, as well as 25565/TCP, 27015/TCP, and 2272/TCP, which are not normally used.

The 53/UDP backscatter that has been on the increase since February this year continued to climb in the current survey period, becoming the port with the highest number of packets observed. Most of these packets show characteristics of the attack method known as DNS Water Torture^{*32}. Additionally, the senders' addresses for the packets observed cover a wide range.

Regarding particularly large numbers of backscatter packets observed by port, there were attacks on the Web servers (80/TCP) of a television station in Ukraine between July 23 and August 15, and on the servers of a hosting provider in Russia between July 1 and July 24. In the latter case in particular, attacks had been continuing since June 16, in the last survey period. In September, many attacks on 25565/TCP were observed. This port is sometimes used with the servers for certain games. These attacks covered a wide range of targets, including those observed on a number of servers of a hosting provider in Russia. Between September 9 and September 10, attacks were observed on DNS (53/TCP) targeting a number of DNS servers that cover the .pk zone for Pakistan domains.

Notable DDoS attacks during the current survey period that were detected via IIJ's observations of backscatter included those between July and early August on sites related to the Israeli government thought to have been carried out by Anonymous.

1.3.2 Malware Activities

Here, we will discuss the results of the observations of the MITF^{*33}, a malware activity observation project operated by IIJ. The MITF uses honeypots^{*34} connected to the Internet in a manner similar to general users in order to observe communications arriving over the Internet. Most appear to be communications by malware selecting a target at random, or scans attempting to locate a target for attack.

*30 Honeypots established by the MITF, a malware activity observation project operated by IIJ. See also "1.3.2 Malware Activities."

*31 The mechanism and limitations of this observation method, as well as some of the results of IIJ's observations, are presented in Vol.8 on this report (http://www.ijj.ad.jp/en/company/development/iir/pdf/iir_vol08_EN.pdf) under "1.4.2 Observations on Backscatter Caused by DDoS Attacks."

*32 Secure64 Software Corporation, "Water Torture: A Slow Drip DNS DDoS Attack" (<https://blog.secure64.com/?p=377>). For an explanation in Japanese, see the following document written by Mr. Morishita of Japan Registry Services. "DNS Water Torture Attacks" (http://2014.secon.jp/dns/dns_water_torture.pdf) (in Japanese).

*33 An abbreviation of Malware Investigation Task Force. The Malware Investigation Task Force (MITF) began activities in May 2007, observing malware network activity through the use of honeypots in an attempt to understand the state of malware activities, to gather technical information for countermeasures, and to link these findings to actual countermeasures.

*34 A system designed to simulate damages from attacks by emulating vulnerabilities, recording the behavior of attackers, and the activities of malware.

■ Status of Random Communications

Figure 5 shows the distribution of sender's IP addresses by country for communications coming into the honeypots between June 1 and September 30, 2014. Figure 6 shows trends in the total volumes (incoming packets). The MITF has set up numerous honeypots for the purpose of observation. We have taken the average per honeypot, showing the trends for incoming packet types (top ten) over the entire period subject to study. Additionally, in these observations we corrected data to count multiple TCP connections as a single attack when the attack involved multiple connections to a specific port, such as attacks on MSRPC.

Much of the communications arriving at the honeypots demonstrated scanning behavior targeting TCP ports utilized by Microsoft operating systems. We also observed scanning behavior targeting 1433/TCP used by Microsoft's SQL Server, 22/TCP used for SSH, 53/UDP used for DNS, 23/TCP used for Telnet, and 8080/TCP used for HTTP proxies.

During the current survey period, a large volume of 53/UDP communications occurred on July 27, August 14, and September 12. Most of these packets were caused by the receipt of DNS water torture packets, as mentioned in our backscatter observations^{*35}. A small number of queries ranging from one to several were sent to each honeypot from a large number of IP addresses allocated mainly to the United States, Canada, and China. From this trend, we estimate that the attacker used a method such as a botnet. The queries contained A record resolution requests for "(random string).(existing domain)." On September 8, communications targeting 3395/UDP were made to the IP address of a specific honeypot from an IP address allocated to Iran. Upon investigating these communications, we found that random data from several dozen to several hundred bytes in length had been sent.

■ Malware Network Activity

Figure 7 shows the distribution of the specimen acquisition source for malware during the period under study, while Figure 8 shows trends in the total number of malware specimens acquired. Figure 9 shows trends in the number of unique specimens.

In Figure 8 and Figure 9, the number of acquired specimens show the total number of specimens acquired per day^{*36}, while the number of unique specimens is the number of specimen variants categorized according to their digest of a hash function^{*37}. Specimens are also identified using anti-virus software, and a breakdown of the top 10 variants is displayed color coded by malware name. As with our previous report, for Figure 8 and Figure 9 we have detected Conficker using multiple anti-virus software packages, and removed any Conficker results when totaling data.

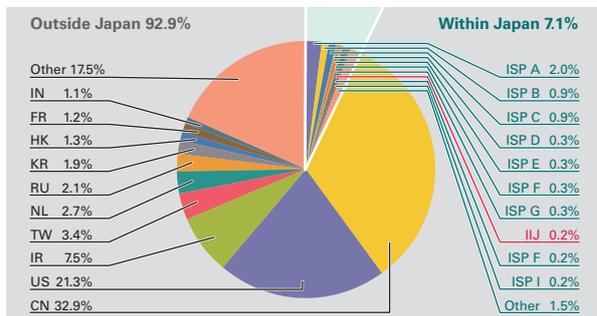


Figure 5: Sender Distribution (by Country, Entire Period under Study)

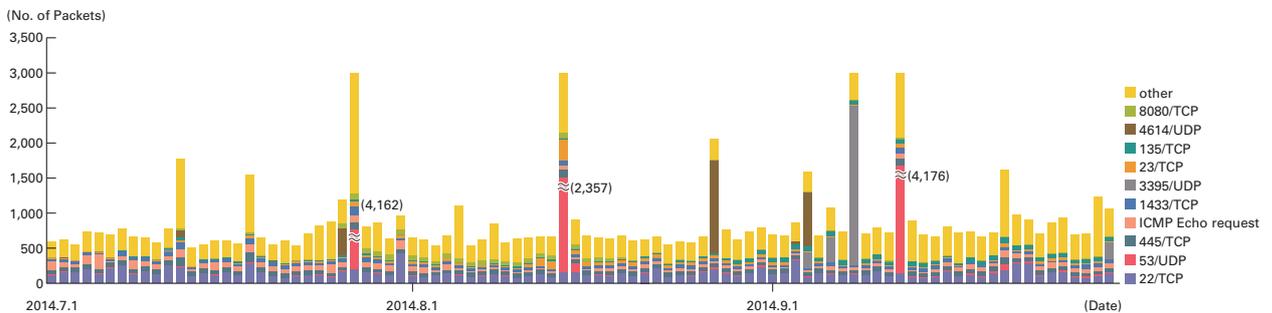


Figure 6: Communications Arriving at Honeypots (by Date, by Target Port, per Honeypot)

*35 The MITF honeypots do not query authoritative servers or cache servers when they receive DNS query packets, so they provide no aid to attacks.

*36 This indicates the malware acquired by honeypots.

*37 This figure is derived by utilizing a one-way function (hash function) that outputs a fixed-length value for various input. The hash function is designed to produce as many different outputs as possible for different inputs. While we cannot guarantee the uniqueness of specimens by hash value, given that obfuscation and padding may result in specimens of the same malware having different hash values, the MITF has expended its best efforts to take this fact into consideration when using this methodology as a measurement index.

On average, 93 specimens were acquired per day during the period under study, representing 20 different malware. After investigating undetected specimens more closely, malware that steals passwords was observed from IP addresses allocated to Singapore and the Philippines. Additionally, about 54% of undetected specimens were in text format. Because many of these text format specimens were HTML 404 or 403 error responses from Web servers, we believe this was due to infection behavior of malware such as old worms continuing despite the closure of download sites that newly-infected PCs access to download malware.

Under the MITF's independent analysis, during the current period under observation 96.0% of malware specimens acquired were worms, and 4.0% were downloaders. In addition, the MITF confirmed the presence of 1 botnet C&C server*38 and 16 malware distribution sites.

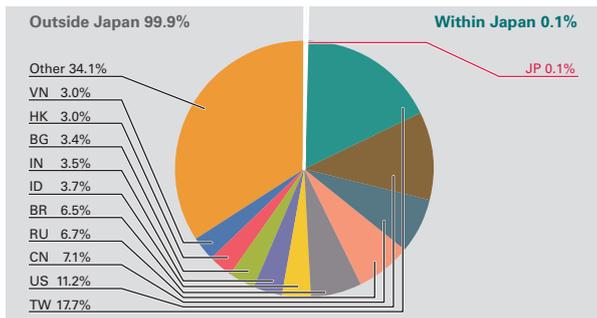


Figure 7: Distribution of Acquired Specimens by Source (by Country, Entire Period under Study, Excluding Conficker)

■ Conficker Activity

Including Conficker, an average of 25,435 specimens were acquired per day during the period covered by this report, representing 672 different malware. While figures rise and fall over short periods, Conficker accounts for 99.6% of the total number of specimens acquired, and 97.0% of unique specimens. This demonstrates that Conficker remains the most prevalent malware by far, so we have omitted it from figures in this report. The total number of specimens acquired during the period covered by this report decreased by approximately 20% compared to the previous survey period. Unique specimens were also down

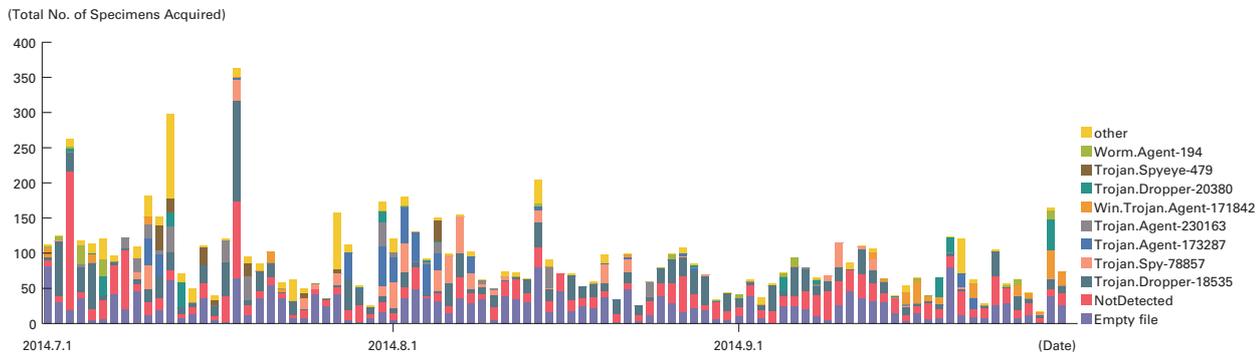


Figure 8: Trends in the Total Number of Malware Specimens Acquired (Excluding Conficker)

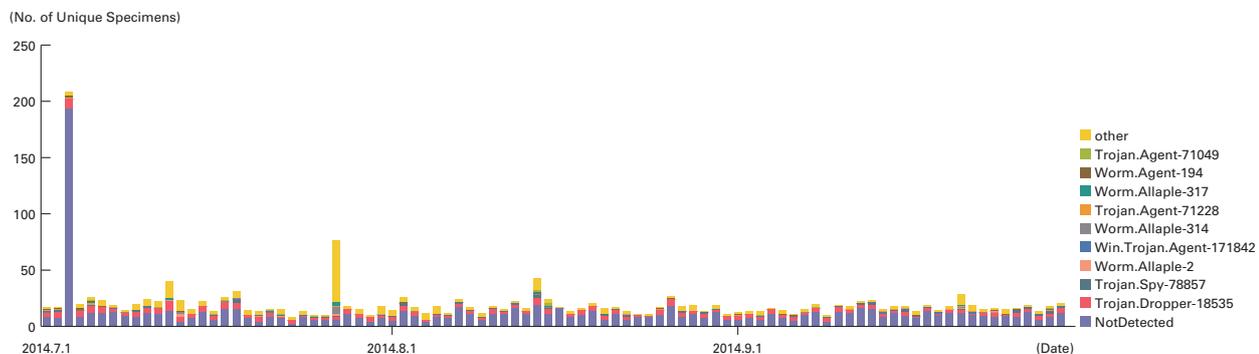


Figure 9: Trends in the Number of Unique Specimens (Excluding Conficker)

*38 An abbreviation of Command & Control Server. A server that provides commands to a botnet consisting of a large number of bots.

by about 7%. According to the observations of the Conficker Working Group*39, as of September 30, 2014, a total of 1,026,417 unique IP addresses are infected. This indicates a drop to about 32% of the 3.2 million PCs observed in November 2011, but it demonstrates that infections are still widespread.

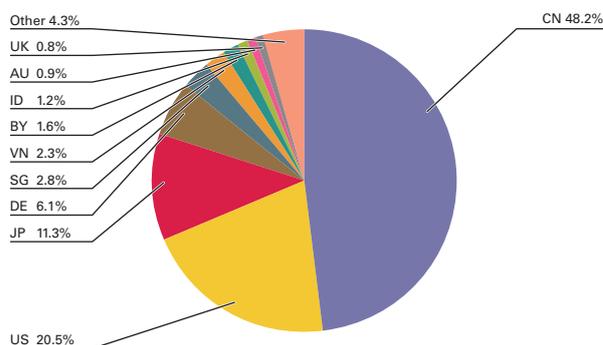
1.3.3 SQL Injection Attacks

Of the types of different Web server attacks, IJ conducts ongoing surveys related to SQL injection attacks*40. SQL injection attacks have flared up in frequency numerous times in the past, and remain a major topic in Internet security. SQL injections are known to occur in one of three attack patterns: those that attempt to steal data, those that attempt to overload database servers, and those that attempt to rewrite Web content.

Figure 10 shows the distribution of SQL injection attacks against Web servers detected between July 1 and September 30, 2014. Figure 11 shows trends in the numbers of attacks. These are a summary of attacks detected by signatures on the IJ Managed IPS Service.

China was the source for 48.2% of attacks observed, while the United States and Japan accounted for 20.5% and 11.3%, respectively, with other countries following in order. There was a dramatic increase in the number of SQL injection attacks against Web servers compared to the previous report. This was due to a significant spike in attacks originating from China.

During this period, attacks from a specific attack source in China directed at specific targets took place on July 19. On August 25, attacks were made from a number of sources in the United States directed at specific targets. On September 8, there were large-scale attacks from specific attack sources in China directed at specific targets. Attacks from a specific attack source in China directed at specific targets also took place on September 23. These attacks are thought to have been attempts to find vulnerabilities on Web servers.



As previously shown, attacks of various types were properly detected and dealt with in the course of service. However, attack attempts continue, requiring ongoing attention.

Figure 10: Distribution of SQL Injection Attacks by Source

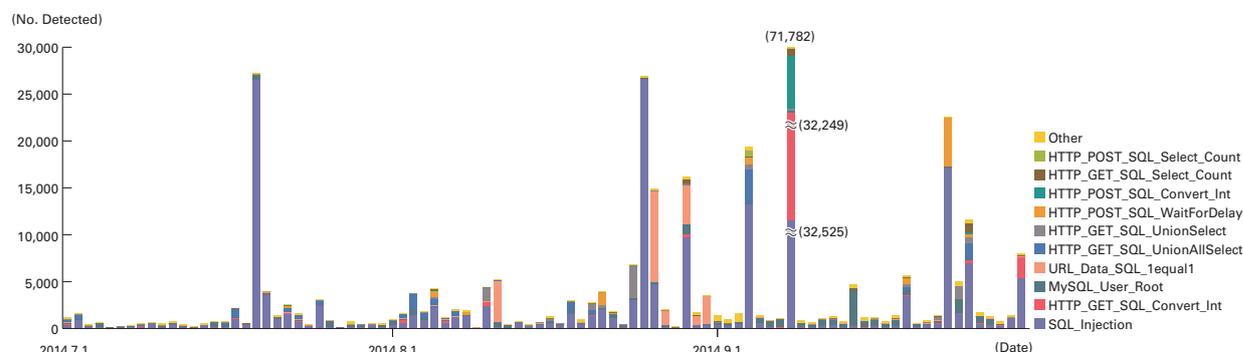


Figure 11: Trends in SQL Injection Attacks (by Day, by Attack Type)

*39 Conficker Working Group Observations (<http://www.confickerworkinggroup.org/wiki/pmwiki.php/ANY/InfectionTracking>).

*40 Attacks accessing a Web server to send SQL commands, thereby manipulating an underlying database. Attackers access or alter the database content without proper authorization, and steal sensitive information or rewrite Web content.

1.3.4 Website Alterations

Here we indicate the status of website alterations as surveyed through the MITF Web crawler (client honeypot)*⁴¹. This Web crawler accesses tens of thousands of websites on a daily basis, with a focus on well-known and popular sites in Japan. We also add new target sites on a regular basis. In addition to this, we temporarily monitor websites that have seen short-term increases in access numbers. By surveying websites thought to be viewed frequently by typical users in Japan, it is easier to speculate on trends regarding fluctuations in the number of altered sites, as well as the vulnerabilities exploited and malware distributed.

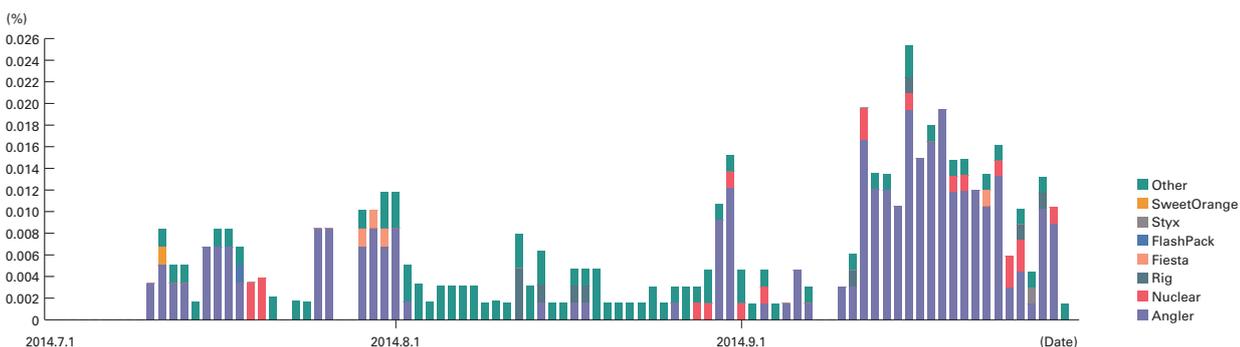
In the period from July to September, 2014, drive-by downloads were observed about twice as often as for the April to June period (Figure 12). As with the previous survey period, most of the attacks were based on Angler and Nuclear. Each of these exploit kits have functions for exploiting vulnerabilities in plug-ins such as Java, Flash, and Silverlight. In particular, functions have been added to Angler at a dizzying pace, and in late August a system that attempts to avoid detection by anti-virus software by not saving malware files on the client's hard disk during infection was discovered*⁴².

Also, after the attack detection logic of our Web crawler system was partially improved in mid-September, the number of Angler cases observed rose considerably. From this, we can surmise that although Angler appeared to have been wiped out in August, it is highly likely that attacks using the same exploit kit actually continued.

Regarding new trends, from around mid-August Rig began to be observed. This is a comparatively new exploit kit that, like the previously-mentioned Angler and Nuclear, has functions for attacking vulnerabilities in Java, Flash, and Silverlight. It is provided as a paid service*⁴³. In some of the cases in which Rig-based attacks were observed, redirection from a single redirection website to multiple exploit kits was seen.

Meanwhile, in many cases where the websites altered to redirect users were not well known, they stayed altered for two to four weeks at a time, and examples in which no investigation of the cause of alteration or fundamental countermeasures were carried out remain rife.

We estimate that drive-by download incidence as a proportion of the whole is in an upward trend. It is recommended that all parties affected continue to exercise caution. Website operators should ensure that measures against the alteration of web content are in place, and visitors should stay up to date with measures against vulnerabilities in browsers or related plug-ins.



*Covers several tens of thousands of sites in Japan. In recent years, drive-by downloads have been configured to change attack details and whether or not attacks are made based on the client system environment or session information, source address attributes, and the quota achievement status of factors such as number of attacks. This means that results can vary wildly at times depending on the test environment and circumstances.

*Because the Web crawler was not operating between July 1 and July 7, no attacks were detected during that period.

Figure 12: Rate of Drive-By Download Incidence When Viewing Websites (%) (by Exploit Kit)

*⁴¹ See "1.4.3 Website Defacement Surveys Using Web Crawlers" in Vol.22 of this report (http://www.iij.ad.jp/en/company/development/iir/pdf/iir_vol22_EN.pdf) for an explanation of Web crawler observation methods.

*⁴² See Malware don't need Coffee, "Angler EK : now capable of 'fileless' infection (memory malware)" (<http://malware.dontneedcoffee.com/2014/08/angler-ek-now-capable-of-fileless.html>) for more information about these functions.

*⁴³ See Kahu Security, "RIG Exploit Pack" (<http://www.kahusecurity.com/2014/rig-exploit-pack/>) for more information about Rig.

1.4 Focused Research

Incidents occurring over the Internet change in type and scope from one minute to the next. Accordingly, IJ works toward implementing countermeasures by continuing to perform independent surveys and analyses of prevalent incidents. Here we will present information from the surveys we have undertaken during this period regarding the Shellshock Bash vulnerability, as well as the POODLE attack. We will also cover the status of list-based attacks and their countermeasures.

1.4.1 The Shellshock Bash Vulnerability

■ About Shellshock

The CVE-2014-6271 vulnerability^{*44} was disclosed along with a fixed version of Bash^{*45} on September 24, 2014. This vulnerability could allow remote arbitrary code execution. Because Bash is a shell program that is normally used locally, people tend to think it is unaffected by remote attacks, but as the impact became clear there was a big uproar.

Although a fixed version was released at the same time that the vulnerability was disclosed, during the same week later several new related vulnerabilities were reported in quick succession. Table 1 shows a list of the related vulnerabilities. CVE-2014-7169^{*46} was registered as a new vulnerability because the original fix was deemed inadequate. CVE-2014-7186^{*47} and CVE-2014-7187^{*48} were vulnerabilities discovered during the process of making fixes. CVE-2014-6277^{*49} and CVE-2014-6278^{*50} were produced during the process of upstream maintainers merging the patches. To prevent attacks based on Shellshock, all these vulnerabilities must be patched.

■ Dealing with the Vulnerabilities

Because Bash is a shell program, unlike software such as server programs or scripting language interpreters, the functionality of the latest version is not often required. This means that the binary package provided for each distribution is used in most cases. Shellshock involves a number of vulnerabilities, but because some of these only affect fixes merged into upstream, the vulnerabilities that require action differ depending on the patches applied.

As an example for one distribution, Table 2 shows the status of each vulnerability in each release of CentOS^{*51}. The version that fixes CVE-2014-6271 contains inadequate fixes just like the others, but subsequent vulnerabilities were fixed at the same time, and it is not affected by the CVE-2014-6277 and CVE-2014-6278 vulnerabilities that result from upstream fixes.

Table 1: List of Shellshock-Related Vulnerabilities

CVE ID	Date of Disclosure	Impact	Notes
CVE-2014-6271	September 24, 2014	Arbitrary code execution	The original vulnerability
CVE-2014-7169	September 25, 2014	Arbitrary code execution	Resulted from inadequate fixes to CVE-2014-6271
CVE-2014-7186	September 26, 2014	Arbitrary code execution	-
CVE-2014-7187	September 26, 2014	DoS	-
CVE-2014-6277	September 27, 2014	DoS	Resulted from upstream version fixes
CVE-2014-6278	September 27, 2014	Arbitrary code execution	Resulted from upstream version fixes

Table 2: Fixed Bash Packages for Each CentOS Release

CVE ID	CentOS 5	CentOS 6	CentOS 7
CVE-2014-6271	bash-3.2-33.el5.1	bash-4.1.2-15.el6_5.1	bash-4.2.45-5.el7_0.2
CVE-2014-7169			
CVE-2014-7186	bash-3.2-33.el5_10.4	bash-4.1.2-15.el6_5.2	bash-4.2.45-5.el7_0.4
CVE-2014-7187			
CVE-2014-6277	No fix necessary	No fix necessary	No fix necessary
CVE-2014-6278	No fix necessary	No fix necessary	No fix necessary

*44 CVE-2014-6271 bash: specially-crafted environment variables can be used to inject shell commands (https://bugzilla.redhat.com/show_bug.cgi?id=CVE-2014-6271).

*45 GNU Bash: The GNU Bourne-Again SHell (<http://www.gnu.org/software/bash/>).

*46 CVE-2014-7169 bash: code execution via specially-crafted environment Incomplete fix for CVE-2014-6271 (https://bugzilla.redhat.com/show_bug.cgi?id=CVE-2014-7169).

*47 CVE-2014-7186 bash: parser can allow out-of-bounds memory access while handling redir_stack (https://bugzilla.redhat.com/show_bug.cgi?id=CVE-2014-7186).

*48 CVE-2014-7187 bash: off-by-one error in deeply nested flow control constructs (https://bugzilla.redhat.com/show_bug.cgi?id=CVE-2014-7187).

*49 CVE-2014-6277 bash: uninitialized here document closing delimiter pointer use (https://bugzilla.redhat.com/show_bug.cgi?id=CVE-2014-6277).

*50 CVE-2014-6278 bash: incorrect parsing of function definitions with nested command substitutions (https://bugzilla.redhat.com/show_bug.cgi?id=CVE-2014-6278).

*51 CentOS Project (<http://www.centos.org/>).

This is because the Red Hat Enterprise Linux^{*52} that CentOS is based on dealt with the vulnerabilities by applying patches that fixed them to the base package rather than updating to the latest version of Bash. Whether or not other distributions are affected also depends on how fixes were applied.

■ The Impact of the Vulnerabilities

Let us examine why the vulnerabilities in Bash, which is used as a local shell, turned out to be exploitable remotely. These vulnerabilities apply when a Bash program is launched with specific environment variables set. In some environments Bash is run as the `/bin/sh` program used with shell-related calls, and these are also affected similarly.

Attacks that actually use these vulnerabilities have already been observed, with the most targeted area being CGI^{*53}. To give an easily understandable example of the impact, there are some CGI programs written in Bash that run on Web servers, but these are almost never run on actual environments. However, even if a CGI program is written in another language, in some cases it is affected because it depends on the shell internally.

Linux is also often used as a control OS for appliance products or embedded devices. Some familiar examples are load balancers, home routers, and file servers. These may also have Bash embedded, or a management screen in which CGI is used, meaning that they are affected in the same way as normal Unix servers. In light of this, we will explain program execution in Unix environments, as well as the CGI processing flow, to show the impact of these vulnerabilities.

■ Program Execution in Unix Environments

Unlike Windows environments, in Unix environments it is not possible to create completely new processes. Instead, the procedure involves copying (forking) a parent process, and changing the child process into the process to execute (`exec`). Ultimately, the child process turns into the program to execute, becoming something different to the parent process. However, because it is copied, some of the data is carried over. Figure 13 shows the processing flow.

As shown in Table 3, there exist multiple varieties of `exec` functions called when making this change exist, and they differ in their handling of their handling of the environment variables that are the key factor for these vulnerabilities. For `exec` functions that do not explicitly specify an environment variable when called, the environment variables carried over from the parent process to the child process are still used after it changes into the process to execute.

■ CGI Processing Flow

CGI is a system for providing dynamic content via Web servers. It is less prevalent now that Web application frameworks have become more widespread, but it is still used. Web servers convert the headers received via HTTP into environment variables to launch the CGI program. This procedure is why CGI is the area most easily affected by these vulnerabilities. As long as Bash is not launched directly as a CGI program, it is not subject to the vulnerabilities at this point. However, as

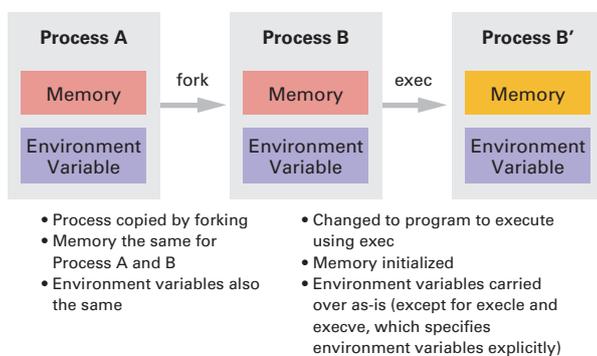


Figure 13: Execution of External Programs in Unix Environments

Table 3: Operation of `exec` Functions

Function Name	Environment Variable Handling
<code>execl</code>	Carried over
<code>execle</code>	Explicitly specified
<code>execlp</code>	Carried over
<code>execv</code>	Carried over
<code>execve</code>	Explicitly specified
<code>execvp</code>	Carried over

*52 Red Hat Enterprise Linux (<http://www.redhat.com/en/technologies/linux-platforms/enterprise-linux>).

*53 RFC3875 The Common Gateway Interface Version 1.1 (<http://www.ietf.org/rfc/rfc3875.txt>).

mentioned earlier, the environment variables are carried over to the child process, so the vulnerabilities may apply if an external program is called via the shell, even if written in another language. In addition to directly-called programs, the same goes for programs called indirectly via libraries or modules. Figure 14 shows whether or not the vulnerabilities apply when programs are called via CGI.

With scripting language the shell may also be called unintentionally depending on how a program is written. Table 4 shows operation in Perl^{*54} as an example of scripting language. As this demonstrates, even calls via the same function may be implicitly expanded to calls via the shell based on the parameters. Because all environment variables are carried over, when they are expanded to calls via the shell, the Shellshock vulnerabilities apply.

■ Summary

As explained above, even vulnerabilities in a program like Bash, which appears at first to be untouchable by remote attacks, may be attacked by external parties depending on how the program is used.

With these kinds of criteria for determining impact, it is difficult to comprehensively investigate all execution paths via a static inspection based on the source code. Meanwhile, unlike server programs, in many cases Bash updates do not require a reboot, so we recommend updating to a fixed version quickly if there is even the slightest suspicion that a system is affected.

Appliance products or embedded devices will require a firmware patch from the vendor. For some products it will take time before fixed firmware is released, and in this case provisional measures such as restricting access to areas suspected of being affected are also effective.

There are a range of techniques for dealing with vulnerabilities, such as applying workarounds, updating to a fixed version, or implementing protection using security devices. The best measure depends on the configuration of the system in question, so it is important to deal with the issue appropriately by evaluating each particular situation.

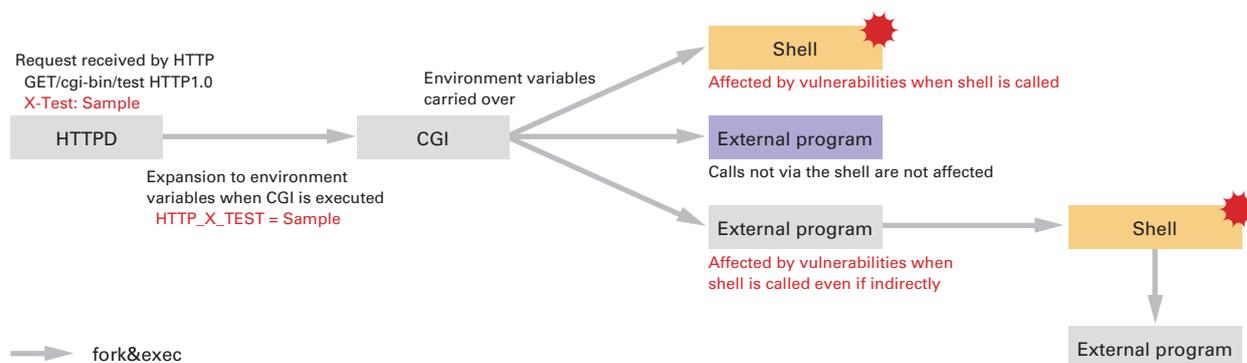


Figure 14: Impact on CGI

Table 4: External Program Calls in Perl

Perl Code	Expansion to exec Function	Shell Call	Environment Variables
<code>"/bin/ls"</code>	<code>execvp("/bin/ls", ["/bin/ls"])</code>	No	Carried over
<code>"/bin/ls 2>/dev/null"</code>	<code>execl("/bin/sh", ["-c", "/bin/ls 2> /dev/null"])</code>	Yes	Carried over
<code>system("/bin/ls")</code>	<code>execvp("/bin/ls", ["/bin/ls"])</code>	No	Carried over
<code>system("/bin/ls &")</code>	<code>execl("/bin/sh", ["-c", "/bin/ls &"])</code>	Yes	Carried over
<code>open(FH, "/bin/ls ")</code>	<code>execvp("/bin/ls", ["/bin/ls"])</code>	No	Carried over
<code>open(FH, "/bin/ls 2>/dev/null ")</code>	<code>execl("/bin/sh", ["-c", "/bin/ls 2>/dev/null "])</code>	Yes	Carried over

*54 The Perl Programming Language (<https://www.perl.org/>).

1.4.2 The POODLE Attack

On October 14, 2014 U.S. time, a new attack on SSLv3^{*55} was disclosed by Google's research team^{*56}. This technique, known as the POODLE attack (Padding Oracle On Downgraded Legacy Encryption attack), is a MITM attack similar to the BEAST attack^{*57}. It is used to decrypt target data encrypted with SSL one byte at a time through trial and error by sending large numbers of requests from a browser to a server. As far as practical damages go, it could result in the theft of cookies. Additionally, while it may appear at first as if this technique could also be applied to TLSv1.0^{*58}, because of the differences in the padding methods used by SSLv3.0 and TLSv1.0, it is not practical as an attack on TLSv1.0. On the other hand, attacks can be made on SSLv3.0 by simply sending up to 255 queries to a server to sneak a glimpse of one byte. In this section, we examine the technical background of the POODLE attack, as well as its practicality^{*59}.

■ An Overview of CBC Mode and its Potential Vulnerabilities

Like the BEAST attack, the POODLE attack only succeeds when CBC (Cipher Block Chaining mode) is used as the cipher mode. Stream cipher algorithms such as RC4 involve combining a keystream generated based on the key with plaintext using the XOR operation, enabling data of an arbitrary length to be encrypted. Meanwhile, block cipher algorithms such as DES or AES involve inputting data of a fixed length (the block length is 64 bits for DES and 128 bits for AES) into the cipher algorithm to obtain ciphertext of the same block length as the input length. When using block ciphers, the data to be encrypted is normally far larger than the block length, so it is necessary to encrypt many segments sequentially using the block cipher algorithm. The system that determines how the data obtained in the previous block cipher processing is used in the next data processing is called the cipher mode, and a number of systems have been proposed. Figure 15 shows an overview of one of these, the CBC cipher mode^{*60}. First, an IV (Initial Vector)^{*61} of the block length is prepared as the initial value. Actual block cipher processing is carried out on data consisting of P_1 , the first part of the plaintext after it is partitioned into block length segments, combined with the IV using the XOR operation. The output of this is the ciphertext C_1 . To encrypt P_2 , the second plaintext part, it is combined with C_1 obtained from the previous block cipher processing using the XOR operation, and this data is input as plaintext and encrypted, producing C_2 in a similar manner to before. In this way, encryption of the entire plaintext is calculated by repeating the cipher processing of $C_i = \text{ENC}(P_i \oplus C_{i-1})$. Conversely, for decryption the $P_i = \text{DEC}(C_i) \oplus C_{i-1}$ decryption processing is repeated sequentially to restore the plaintext.

For this cipher processing, when handling data that cannot be segmented at exactly block length, it becomes necessary to perform padding processing. This involves padding the last part of the plaintext with some form of data when it does not match the block length. Once padded to match the block length, block cipher processing can be carried out.

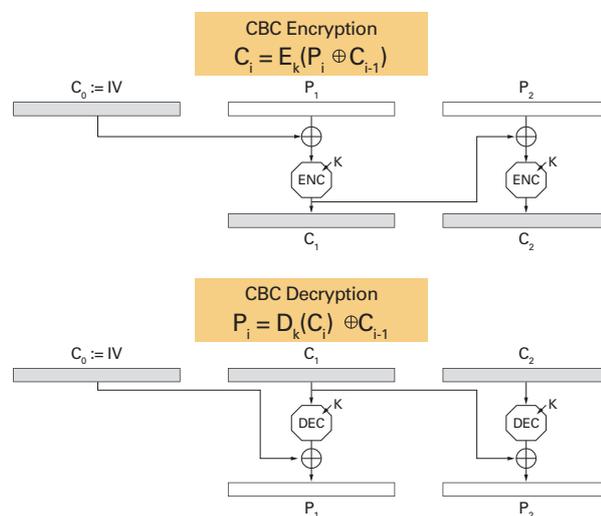


Figure 15: CBC Cipher Mode Overview

*55 IETF, "The SSL Protocol Version 3.0" (<http://tools.ietf.org/html/rfc6101>).

*56 Bodo Möller, Thai Duong, Krzysztof Kotowicz, "This POODLE Bites: Exploiting The SSL 3.0 Fallback" (<https://www.openssl.org/~bodo/ssl-poodle.pdf>).

*57 Thai Duong, Juliano Rizzo, "BEAST - Here Come The XOR Ninjas", 2011.

*58 IETF, "The TLS Protocol Version 1.0" (<http://tools.ietf.org/html/rfc2246>).

*59 This section places emphasis on the technical explanation. For an overview of the POODLE attack, see the following report from the Cryptographic protocol Evaluation toward Long-Lived Outstanding Security Consortium. On a new attack on SSLv3 specification called POODLE attack (https://www.cellos-consortium.org/index.php?PoodleAttack_20141015).

*60 NIST, "NIST Special Publication 800-38A Recommendation for Block Cipher Modes of Operation: Methods and Techniques" (<http://csrc.nist.gov/publications/nistpubs/800-38a/sp800-38a.pdf>).

*61 The IV (Initial Vector) does not need to be kept secret. However, as with a key, the IV is required at the time of decryption.

TLS uses the PKCS#5 padding scheme^{*62}, and Figure 16 shows an example of padding when the block length is 8 bytes (for example DES). Because case (a) is 3 bytes short of the block length, that 3 byte space is filled with 0x03 as padding data. In the same way, case (b) is 5 bytes short so that 5 byte space is filled with 0x05. Although (c) is a case in which the data to be encrypted is a multiple of the block length, to identify the boundary showing how much padding was added to the plaintext, it is treated as 8 bytes short (exactly the block length), and this block is completely filled with 0x08.

Meanwhile, SSLv3.0 uses a different padding scheme, and this difference is one of the factors that brought about the POODLE attack. Figure 17 shows the differences between PKCS#5 padding and SSLv3 padding. The padding scheme used with SSLv3 assigns the padded byte length to the final byte when padding, just as with PKCS#5 padding. Random data is used for the rest of the padding data. Unfortunately, for both SSLv3 and TLS this padded area does not fall within the scope of the MAC (Message Authentication Code) data that indicates data has not been altered over the communication channel. That means even if the padding data is altered, these alterations cannot be detected via MAC verification. On one hand, while with TLS the constraints of the padding data explained above may enable alterations to the padded part to be detected, with SSLv3 there is no way to detect alterations to the padding data other than the final byte. You could say that the POODLE attack targets this tiny gap in the protocol's armor.

Figure 18 shows the extent of impact for decryption using CBC mode when the ciphertext has been altered. In this figure, the entire P_1 plaintext block has been changed by altering part of the C_1 ciphertext. However, the attacker is able to freely change the intended area of the P_2 plaintext. Figure 19 shows the principles of padding oracle attacks^{*63} that exploit this characteristic. Assuming that the plaintext length is known, and data up to P_n is encrypted, the man-in-the-middle replaces the last block C_n with the attack target, or in other words the C_i data they want to decrypt. Because the plaintext length is already known at this time, the last byte of P_n that stores the padded data length is also known. SSL/TLS servers conduct

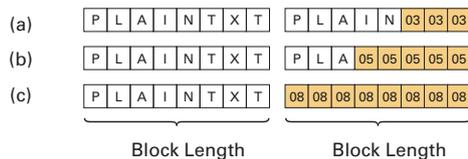
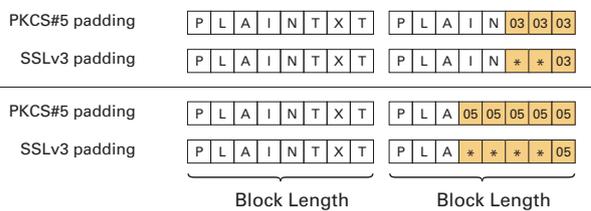
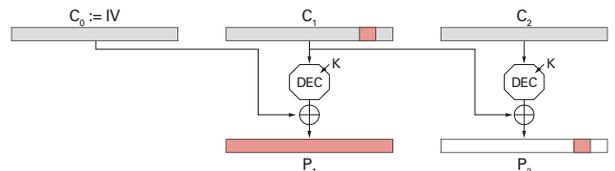


Figure 16: PKCS#5 Padding Example with Block Length of 8 Bytes



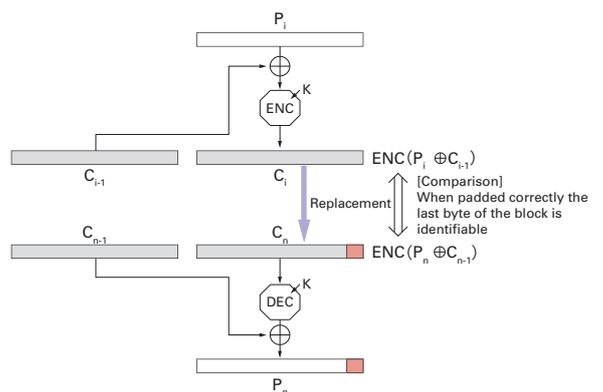
The "" marks in colored padding data indicate random data.

Figure 17: Differences Between PKCS#5 Padding Scheme and SSLv3 Padding Scheme



*When a single byte of the ciphertext is altered and decrypted, this affects the entire block of the plaintext for the corresponding block, but it allows the intended area of the plaintext for the next block to be altered.

Figure 18: Extent of Impact When Ciphertext is Altered in CBC Mode



*The server's padding check function accepts data when padding is correct, but returns an error when it is incorrect.

Figure 19: Padding Oracle Attack Principles

^{*62} IETF, "PKCS #5: Password-Based Cryptography Specification Version 2.0" (<http://tools.ietf.org/html/rfc2898>).

^{*63} In addition to SSL/TLS, the same kind of attack has been disclosed for SSH and IPsec. In both cases, these attack methods decipher encrypted data by determining whether it conforms to data format constraints determined by the specifications based on error information from the server. M.R. Albrecht, K.G. Paterson, G.J. Watson, Plaintext Recovery Attacks Against SSH, 30th IEEE Symposium on Security and Privacy, 2009. (<http://www.isg.rhul.ac.uk/~kp/SandPfinal.pdf>). J.P. Degabriele and K.G. Paterson, On the (In)security of IPsec in MAC-then-Encrypt Configurations, Proceedings of the 17th ACM Conference on Computer and Communications Security, 2010. (<http://portal.acm.org/citation.cfm?id=1866363>).

a padding check according to procedure, and an error is returned if the data does not conform to the padding scheme determined by the specifications. Using this error return function (padding oracle) means that when the server accepts the encrypted data, the correct data is stored in the padded part. In SSLv3, the data is accepted if the last byte of P_i and P_n match. As a result, it is possible to decrypt the last byte of P_i (even when the encryption key is not known) because the last byte of $P_i \oplus C_{i-1}$ and $P_n \oplus C_{n-1}$ match. In actual fact, because a MAC-based check is implemented, it is possible to detect whether the plaintext part before the padding data has been altered, so in most cases this doesn't work. However, as pointed out with the POODLE attack, when that block is somehow filled solely with padding data (as with (c) in Figure 16), the final block can be removed from the scope of MAC verification. When testing whether the last byte matches, it will definitely be accepted once every 256 times. Meanwhile, for TLS all bytes in the block must be filled with the same data, as shown in (c) of Figure 16. In this example the block length is 8 bytes, so just one in every 2^{64} attempts will be successful. For AES the block length is 16 bytes, making 2^{128} attempts necessary, meaning for TLS this attack is regarded as a potential problem that is not currently practical. However, the POODLE attack identified here only requires around 2^8 calculations to decrypt 1 byte, so it is recognized as a practical attack.

■ The Practicality of the POODLE attack

For SSH, the CTR cipher mode was available in addition to the CBC mode, so by shelving the use of CBC mode it was possible to prevent padding oracle attacks. However, for SSL/TLS the use of CTR mode is not established in the specifications. Additionally, for the CipherSuites (cryptographic algorithms) specified in the SSLv3 specifications, other than algorithms with only 40-bit key strength that were formerly used for export controls, or NULL (no encryption), the only option available when CBC mode is not used is RC4. However, last year new methods for attacking RC4 were disclosed in quick succession^{*64*65}, and it is now seen as a vulnerable algorithm.

The BEAST attack is classified as a padding oracle attack like the POODLE attack, and can be made successfully on SSLv3 and TLSv1.0. The TLSv1.1 specification is designed to work around this vulnerability. More specifically, it avoids the underlying problem of cipher data from the previous session being used as the IV for the next session by implementing a process that generates a new IV. SSLv3 and TLSv1.0 are still in use despite the BEAST attack existing because the technique called 1/n-1 splitting is known to be an effective countermeasure^{*66}. This technique has already been implemented in major browsers, but it cannot be applied to POODLE attacks. It is possible that techniques like 1/n-1 splitting for prolonging the life of SSLv3 will be devised in the future. However, at this point in time there is no longer any way to use SSLv3 safely.

There are two fundamental countermeasures for this attack: 1) disable SSLv3 (on either the client or server), or 2) implement TLS_FALLBACK_SCSV (both client and server must support it). With regard to the first approach, major browsers have announced they will disable SSLv3 in the future^{*67*68*69}. Additionally, instructions for disabling SSLv3 in browsers yourself have been made available^{*70*71}. Methods for disabling SSLv3 on the server side are also being shared^{*72}. Meanwhile, we believe it will be difficult to implement measures for legacy products such as feature phones and game devices. For that reason, once SSL/TLS servers remove support for SSLv3, in some cases it will no longer be possible to view sites on these devices.

*64 Takatori ISOBE, Toshihiro OHIGASHI, Yuhei WATANABE, and Masakatu MORII, "Full Plaintext Recovery Attack on Broadcast RC4," Proc. the 20th International Workshop on Fast Software Encryption (FSE 2013), 2013. (Revised Selected Papers, LNCS 8424, pp.179-202, Springer-Verlag, 2014).

*65 N.AlFardan, D.J.Bernstein, K.G.Paterson, B.Poettering, J.C.Schuldt, On the Security of RC4 in TLS, USENIX Security 2013. (<http://dl.acm.org/citation.cfm?id=2534793>).

*66 ImperialViolet, "POODLE attacks on SSLv3 (14 Oct 2014)" (<https://www.imperialviolet.org/2014/10/14/poodle.html>).

*67 Mozilla Security Blog, "The POODLE Attack and the End of SSL 3.0" (<https://blog.mozilla.org/security/2014/10/14/the-poodle-attack-and-the-end-of-ssl-3-0/>). *68 Google Online Security Blog, "This POODLE bites: exploiting the SSL 3.0 fallback" (<http://googleonlinesecurity.blogspot.jp/2014/10/this-poodle-bites-exploiting-ssl-30.html>).

*69 "Microsoft Security Advisory 3009008, Vulnerability in SSL 3.0 Could Allow Information Disclosure" (<https://technet.microsoft.com/en-us/library/security/3009008>).

*70 SANS ISC InfoSec Handlers Diary Blog, "OpenSSL: SSLv3 POODLE Vulnerability Official Release" (<https://isc.sans.edu/diary/OpenSSL%3A+SSLv3+POODLE+Vulnerability+Official+Release/18827>).

71 mozillaZine, "Security.tls.version." (http://kb.mozillazine.org/Security.tls.version.*).

*72 F5 Networks, "CVE-2014-3566: Removing SSLv3 from BIG-IP" (<https://devcentral.f5.com/articles/cve-2014-3566-removing-sslv3-from-big-ip>).

The second countermeasure is a method proposed for fundamentally protecting against downgrade attacks^{*73*74}. For this method the TLS_FALLBACK_SCSV CipherSuite value and inappropriate_fallback error code are added, making it possible to protect against downgrade attacks as long as both the client and server support this SCSV (Signaling Cipher Suite Value)^{*75}. Its behavior is extremely simple, with a client including this SCSV in the CipherSuites of the ClientHello message when it connects to a server using a protocol other than the highest version it supports. Meanwhile, a server receiving a CipherSuites with this SCSV returns an error if the protocol version specified by the client is lower than the highest version supported by the server. Through this behavior, it is possible to avoid potential downgrade attacks. Currently the draft is at last call status in the TLS WG, but SCSV has been implemented on Google servers and in Chrome (version 33 or later) since February of this year. The version of OpenSSL released on October 15 also implements SCSV. Support is planned for Firefox^{*76}, and in the future we expect use of SCSV on servers and clients will continue to gather momentum. However, until both servers and clients support this countermeasure, it will not be possible to prevent the POODLE attack. It is very unlikely that these countermeasures will be widely applied in legacy environments such as feature phones, so you could regard these as being implemented to guard against potential attacks in the future.

Because generating a large volume of requests from a browser and rewriting some of them over the communication path (to replace the targeted cipher data) are prerequisites of the POODLE attack, in a similar manner to the BEAST attack it is not necessarily possible to easily launch attacks with typical usage environments. However, in this case the issue is with the protocol specification itself, and due to the difficulty of implementing workarounds, SSLv3 is seen as vulnerable like SSLv2, so we do not believe it is wise to continue using it in the future. In fact, Twitter and a number of other services announced they would disable SSLv3 right away^{*77*78}.

In addition to the TLSv1.0 TLS specification, TLSv1.1 and TLSv1.2 have also been standardized and are widely implemented. The former fixes the BEAST attack, and the latter features the more secure SHA-256 and SHA-384 hash functions, as well as support for GCM and CCM cipher modes besides CBC. Furthermore, TLSv1.3^{*79} is currently under discussion in the TLS WG. The elimination of CBC as a cipher mode is being looked into, and preparation for enabling secure use through using the latest protocol version is well underway. Just as the use of cryptographic algorithms such as DES and MD5, which were formerly considered secure and widely used, ceased when they were compromised^{*80}, it will also be necessary to discard past versions of cryptographic protocols and swiftly migrate to new versions. In doing so, we will need to weigh the balance between backward compatibility (the ratio of SSL clients that will no longer be able to connect after moving to a newer version) and security (the urgency of implementing countermeasures).

*73 MITM attacks that involve intentionally forcing the server and client to use an older protocol version. The ClientHello contains a space for specifying the version that the client supports, and by specifying SSLv2 here it is possible to force the use of a weaker algorithm involuntarily. In SSLv2 there is no fundamental solution for downgrade attacks because this message is not protected by MAC, so SSLv2 is recognized as a vulnerable version.

*74 IETF, "TLS Fallback Signaling Cipher Suite Value (SCSV) for Preventing Protocol Downgrade Attacks" (<https://tools.ietf.org/html/draft-ietf-tls-downgrade-scsv>). The draft that was initially a personal proposal (<https://tools.ietf.org/html/draft-bmoeller-tls-downgrade-scsv-02>) is now taken up for discussion in the TLS WG.

*75 Originally, the CipherSuite was 2-byte data that indicated the class of cryptographic algorithm used in SSL/TLS. However, as with the TLS_EMPTY_RENEGOTIATION_INFO_SCSV value introduced in RFC5746 (<http://tools.ietf.org/html/rfc5746>) to deal with the renegotiation function problem discovered in 2009, it has been expanded to serve different purposes than originally intended.

*76 <http://www.ietf.org/mail-archive/web/tls/current/msg13905.html>

*77 <https://twitter.com/twittersecurity/status/522190947782643712>

*78 CloudFlare blog, "SSLv3 Support Disabled By Default Due to POODLE Vulnerability" (<https://blog.cloudflare.com/ssl3-support-disabled-by-default-due-to-vulnerability/>).

*79 The Transport Layer Security (TLS) Protocol Version 1.3 (<https://tools.ietf.org/html/draft-ietf-tls-tls13>). See <https://github.com/tlswg/tls13-spec> for discussion of this version.

*80 Examples of the compromise of cryptographic algorithms are provided in Vol.8 of this report (http://www.ijj.ad.jp/en/company/development/iir/pdf/iir_vol08_EN.pdf) under "1.4.1 Trends in the Year 2010 Issues on Cryptographic Algorithms".

1.4.3 The Status of List-Based Attacks and Their Countermeasures

■ About List-Based Attacks

Since last year there has been a spate of incidents of unauthorized access on a range of online services using attack methods known as list-based attacks and list-based account hackings. List-based attacks are those that involve attackers making ongoing attempts to log in to a service using a list of login ID and password pairs. The lists used in attacks are created from information leaked from services other than the attack target. Attackers use these lists in attacks because many users reuse that same login ID and password on multiple services*⁸¹.

When an attacker logs in successfully, from the service provider's perspective it is difficult to determine whether or not they are a legitimate user, and this enables the attacker to achieve their purpose before measures are taken. The more services that a user reuses the same login ID and password on, the greater the extent of damages to them could be.

■ Differences Compared to Other Login Attempt Attacks

Some well known examples of login attempt attack methods that target specific login IDs include brute-force attacks, which attempt to use a variety of passwords sequentially, and dictionary attacks, which attempt to use words and phrases commonly used as passwords. List-based attacks are different in that the login ID and password pairs used in attempts are already determined in advance.

■ The Status of List-Based Attacks in 2014

Table 5 summarizes the incidents reported as list-based attacks that occurred between January 1 and September 30, 2014. This shows that practically every month services have been targeted by list-based attacks, resulting in unauthorized logins. There were also nine incidents that resulted in financial damages, including cases in which a credit card was used to purchase large numbers of tickets, which were then converted to cash, as well as cases in which messages impersonating a friend were sent to request victims to buy prepaid cards. In addition to these, there were incidents in which service points were converted into gift vouchers without authorization. In 13 cases the list-based attack was discovered by the service provider, while in 9 cases it was the user that discovered the attack. This demonstrates that in many cases the service provider is not aware that a list-based attack has taken place.

■ List-Based Attack Countermeasures on Online Services

Last of all, we explain the list-based attack countermeasures that can be implemented by service providers and users below*⁸². However, these countermeasures are based on the provision/use of a service, and some areas may not be applicable to internal systems.

■ List-Based Attack Countermeasures for Service Providers

The three main countermeasures for service providers are (1) improvements to the design and implementation of authentication functions, (2) improvements to system monitoring and operation, and (3) providing monitoring and notification functions for user activity. Additionally, when making countermeasures, it is necessary to implement them without placing undue burden on users.

The countermeasures under (1) involve improving login ID and password settings, and providing authentication other than passwords. To prevent the same login ID and password being reused, consider auto-generating login IDs and passwords based on system policy. This makes it possible to prevent the reuse of login IDs and passwords from other services*⁸³. When allowing users to set their own passwords, it is necessary to reject weak passwords*⁸⁴ and make it clear that login IDs and passwords should not be reused.

In addition, you must look into supplementing this with multi-step authentication or multifactor authentication (hardware tokens or IC cards) on the assumption that some users will still reuse passwords from elsewhere. To implement these the system will need to be modified in some cases.

*81 The "Stop reusing passwords! A call to prevent unauthorized login due to password list attacks" (<https://www.jpCERT.or.jp/pr/2014/pr140004.html>) (in Japanese) article published by the IPA and JPCERT/CC states that approximately a quarter (25.4%) of users use the same password at money-related service sites.

*82 The Ministry of Internal Affairs and Communications also published "Measures for dealing with unauthorized logins due to list-based account hacking" (http://www.soumu.go.jp/menu_news/s-news/01ryutsu03_02000063.html) (in Japanese) for service providers.

*83 Because it is technically difficult to prevent users from using the login IDs and passwords generated by a service provider on another service, it is necessary to warn users not to use the same details on other services.

*84 It is necessary to prevent users from setting weak passwords that are easy to guess, such as "password," "123456," the user's name, address, or birthday, or a combination of these.

It is also not possible to say with any certainty that information will not leak from your own system due to vulnerabilities, etc. That means it is necessary to take measures to prevent attackers from using authentication information in list-based attacks once they obtain it, in the event that login IDs or passwords leak from your own service. When saving passwords, instead of saving as plain text or using simple hashing, use hash based password stretching with salt to make it impossible to analyze in a practical amount of time. When you cannot implement this yourself, you can also use authentication functions provided by external organizations. However, you must also take into consideration the risk of outsourced authentication systems being attacked.

The countermeasures under (2) apply before and after an attack occurs. When a list-based attack is carried out, in some cases you will observe symptoms not normally seen, such as a higher load on the system, a larger number of failed login attempts, many login attempts being made with login IDs that do not exist, an abnormal number of login attempts from a specific IP address using different login IDs each time, and access from an IP address block not normally seen.

Look at implementing a system that can notify administrators, automatically block attackers' IP addresses, and lock accounts when symptoms such as these are detected*⁸⁵. Additionally, deleting accounts that have not been used in a long time can prevent damages caused by list-based attacks before they happen.

Table 5: List-Based Attack Incidents in 2014

Disclosure Date	Service Type	Service Provider	Period	Unauthorized Login Attempts	Successful Unauthorized Logins	Discovered By	Financial Damages	Main Countermeasures
Jan. 24	Communications-related service	Company A	Jan. 16	Not known	165	Not known	None	(1) (2) (3) (5) (13)
Feb. 24	Blog service	Company B	Not known	Not known	Not known	Service provider	Unauthorized point conversion	(6) (16)
Feb. 25	SNS	Company C	Feb. 5 to Feb. 11	Not known	370	User	None	(1) (6)
Feb. 28	SNS	Company C	Feb. 28	Not known	16,972	Not known	None	(1) (5) (7) (13) (18)
Feb. 28	Communications-related service	Company D	Feb. 24 to Feb. 25	Approx. 15,000	344	Service provider	Unauthorized content purchase	(2) (3) (4) (5) (7) (10) (13)
Mar. 16	Point-related service	Company E	Mar. 16 to Mar. 17	Approx. 940,000	Approx. 19,000	Service provider	None	(4) (5) (7) (9) (13)
Apr. 23	Member site	Company F	Mar. 23 to Apr. 21	4,600,000	78,361	Service provider	None	(3) (5) (7) (8) (9) (13) (14)
Apr. 30	Communications-related service	Company D	Apr. 14 to Apr. 28	Not known	724	Service provider	Unauthorized content purchase	(2) (3) (4) (5) (7) (10) (13)
May 2	Point-related service	Company G	Apr. 19 to Apr. 29	Not known	273	User	Unauthorized point conversion	(2) (4) (7) (9)
May 30	E-commerce service	Company H	Sept. 2013 to Mar. 2014	Not known	Not known	Not known	Unauthorized content purchase	(19)
June 10	Video service	Company I	May 27 to June 4	2,203,590	219,926	User	Unauthorized point conversion	(13) (17)
June 12	SNS	Company J	Not known	Not known	303	User	Prepaid card purchase	(3) (6)
June 17	SNS	Company C	May 30 to June 17	Approx. 4,300,000	263,596	User	None	(1) (3) (6) (13) (15)
June 20	Blog service	Company B	June 16 to June 19	Approx. 1,600,000	2,398	User	None	(12) (13)
June 23	Blog service	Company K	June 19 to June 23	2,293,543	38,280	Service provider	None	(13)
June 26	Survey service	Company L	June 23 to June 24	Not known	Up to 11,502	User	Unauthorized point conversion	(1) (3) (17)
June 30	Game related	Company M	June 28 to June 29	1,796,629	14,399	Service provider	None	(3) (14)
July 1	Ad service	Company N	Irregularly since Aug. 2013	Not known	Not known	Not known	Not known	(7)
July 4	Survey service	Company O	June 25	3,420,000	15,092	Service provider	Unauthorized point conversion	(3) (5) (13)
July 4	Game related	Company P	Not known	Not known	Not known	User	Not known	(5) (8)
Aug. 13	E-commerce service	Company Q	Aug. 7 to Aug. 12	4,220,382	20,957	Not known	None	(13)
Aug. 18	Point-related service	Company E	Aug. 15	Approx. 296,000	756	Service provider	None	(13)
Aug. 27	Game related	Company P	Not known	Not known	Not known	User	Not known	(5) (8) (11)
Sept. 11	Logistics service	Company R	Sept. 10 to Sept. 11	Approx. 11,520,000	Approx. 21,000	Service provider	None	(5) (7) (8) (9) (13)
Sept. 23	E-commerce service	Company Q	Sept. 22 to Sept. 23	18,663	19	Not known	None	(13)
Sept. 26	Logistics service	Company S	Sept. 25 to Sept. 26	Approx. 190,000	10,589	Service provider	None	(1) (5) (7) (8) (9) (13)
Sept. 29	Logistics service	Company T	Not known	Not known	34,161	Service provider	None	(1) (5) (7) (8) (9) (13)
Sept. 30	Communications-related service	Company U	Sept. 27 to Sept. 29	Approx. 2,250,000	6,072	Service provider	None	(1) (5) (7) (8) (9) (11) (14)

(1) Block attacker's IP address

(2) Bolster detection system

(3) Bolster authentication function

(4) Bolster monitoring system

(5) Notify all users that reuse of passwords is prohibited

(6) Request that all users change their passwords

(7) Request that all users change password regularly

(8) Notify all users that previously used passwords are prohibited

(9) Notify all users that easy to guess passwords are prohibited

(10) Notify all users to be wary of phishing sites

(11) Notify all users that two-step authentication is recommended

(12) Force users logged in without authorization to log out

(13) Reset passwords of users logged in without authorization or request them to change their passwords

(14) Lock the account of users logged in without authorization

(15) Lock dormant accounts

(16) Modify system

(17) Shut down part of system

(18) Delete malicious posts

(19) Assist police with investigation

*⁸⁵ Symptoms such as these reflect current attack trends, and in the future these patterns may change. It is necessary to gather information on the latest attack trends, and perform monitoring with these in mind. Additionally, regarding the legal issues surrounding the privacy of communications, such as the identification of symptoms and automatic blocking based on this, see "Section 4: Dealing with spam that exploits SMTP authentication information" in the "Initial Report of the Workshop on the Appropriate Way to Handle Cyber Attacks in the Telecommunications Business" (http://www.soumu.go.jp/main_content/000283608.pdf) (in Japanese), and "(4) Dealing with spam that exploits SMTP authentication information" in the "Guidelines for Dealing with High Volume Communications and Privacy at Telecommunications Carriers" (http://www.jaipa.or.jp/other/mtcs/guideline_v3.pdf) (in Japanese).

When an unauthorized login is detected, consider shutting down the service to prevent the spread of damage and investigate the matter. The password for the account that was logged into without authorization must be changed to prevent further unauthorized login, such as by resetting it. To be on the safe side, also look into changing the passwords for accounts that were not subject to unauthorized login as well.

The countermeasures under (3) are functions for alerting users when important processes are carried out. At times such as when login succeeds or fails, when personal information is viewed or changed, or when purchase settlement processing is complete, notifying users by email will enable them to realize when unauthorized access has occurred. Additionally, providing functions that enable login history or purchase history to be viewed makes it possible to identify the approximate time that unauthorized access takes place.

■ List-Based Attack Countermeasures for Users

When users set their password themselves, the most effective way to do it is by using password management software to generate and manage different passwords for each service. For web-based services, there is also the method of having the web browser remember passwords created with a password generation tool^{*86}. However, to cope with the risks that arise from using Web browsers, always set a master password^{*87}, and use the latest version of the Web browser.

When a service provides two-step authentication, always endeavor to use it. Also enable notifications in the service, and when a notification showing usage you have no knowledge of arrives, change the password as soon as possible, check that no financial damages have been incurred, and contact the service administrator. It is also a good idea to delete your account when you don't intend to use a service anymore.

1.5 Conclusion

This report has provided a summary of security incidents to which IIJ has responded. In this volume we examined the Shellshock Bash vulnerability, and discussed the POODLE attack. We also summarized the current status and countermeasures for list-based attacks. IIJ makes every effort to inform the public about the dangers of Internet usage by identifying and publicizing incidents and associated responses in reports such as this. IIJ will continue striving to provide the necessary countermeasures to allow the safe and secure use of the Internet.

Authors:



Mamoru Saito

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

Hirohide Tsuchiya (1.2 Incident Summary)

Hirohide Tsuchiya, Tadaaki Nagao, Hiroshi Suzuki, Hisao Nashiwa (1.3 Incident Survey)

Tadashi Kobayashi (1.4.1 The Shellshock Bash Vulnerability)

Yuji Suga (1.4.2 The POODLE Attack)

Minoru Kobayashi (1.4.3 The Status of List-Based Attacks and Their Countermeasures)

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

Contributors:

Masahiko Kato, Masafumi Negishi, Takahiro Haruyama, Yasunari Momoi Office of Emergency Response and Clearinghouse for Security Information, Service Operation Division, IIJ

*86 When managing passwords using a Web browser, there is a chance that passwords will be stolen if the computer is infected with malware, but this can be thought of as one practical method of managing a large number of complicated passwords using a familiar tool.

*87 A master password can be set in Firefox. In other major Web browsers (Internet Explorer, Google Chrome, Safari), passwords are automatically encrypted and saved using the security mechanisms of the OS, etc.

The Latest Streaming Technology

Here we discuss video streaming, which is generating an increasing amount of traffic, and also give an overview of the H.265, 4K, and MPEG-DASH standards.

Additionally, we present a case study of the National High School Baseball Championship at Koshien Stadium.

2.1 Large-Scale Content Delivery

In recent years, the amount of traffic generated by video streaming has been steadily rising. Year after year more large distribution sites are also added via the streaming services that IJ provides. The Visual Networking Index (VNI)*¹ also published by Cisco forecasts that video-based traffic will continue growing. According to this index, the total amount of video traffic is expected to account for between 80% and 90% of global consumer traffic by 2018.

As this demonstrates, video streaming is continuing to grow in scale. In fact, to our surprise even live broadcasts of minor events sometimes draw more views than expected. Forecasting streaming traffic is difficult even for those with experience, but in more and more cases the number of hits turns out to be higher than anticipated.

The widespread use of mobile devices is undoubtedly responsible for streaming seeing this kind of popularity. While subscriber numbers for fixed communication networks (fiber optic lines and CATV, etc.) have stagnated, mobile phone penetration surpassed 100% of the population by the end 2011, and stands at 118% as of June 2014. The increase in wireless speeds for mobile devices, and smartphones in particular, also cannot be overlooked. Trial calculations published for the fifth generation (5G) technology currently in development put its download speed at 100 times that of 2010 under ideal conditions. The traffic generated by mobile devices may even exceed that of past fixed communication networks in the future.

This makes mobile devices an attractive segment for content holders. It has become more and more common to see people using mobile devices while traveling on public transport, or during short periods of waiting. Many content holders are considering and implementing delivery to mobile devices to take advantage of this disposable spare time.

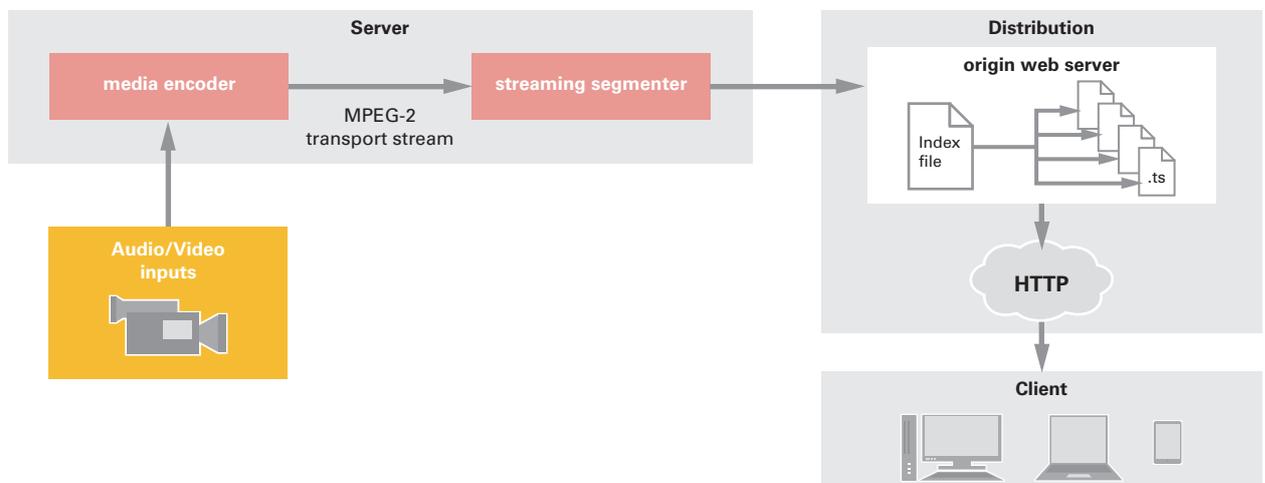


Figure 1: HTTP Live Streaming

*1 Cisco VNI (<http://www.cisco.com/c/en/us/solutions/service-provider/visual-networking-index-vni/index.html>).

2.2 Mobile Topics

A few technical obstacles had to be cleared before delivery to mobile devices was possible. A typical example of this is the streaming formats supported by devices. To deliver content to a client using a format not supported by the OS, the library for that streaming format must be incorporated into an app. Because this method leads to higher licensing and production costs, it is far from ideal. Consequently, it is necessary to address the question of which streaming formats should be pre-installed as standard on mobile OSes.

Apple has supported streaming to its mobile devices since iOS 3.0, which was released in 2009. This was around the time when the iPhone 3GS was released, a year after the iPhone 3G went on sale in Japan. This marked the period when downlink speeds reached a theoretical 7.2 Mbps through support for the HSDPA standard, making streaming to mobile devices more practical. Very few models of the feature phones prevalent at the time were capable of receiving streamed data. Streaming delivery to iPhones only required a Web server, and an increasing number of models were compatible with streaming encoders, so it did not take long for streaming to catch on. People embraced the experience of viewing video content on smartphones.

For Android, the situation differs greatly depending on the OS version. Adobe put out their Flash Player for Android 2.2 and 2.3 when these OSes were released, so most video content was delivered using Flash technology (RTMP: Real Time Messaging Protocol). However, in 2011 Adobe announced they were terminating development of Flash Player for mobile devices. Because HLS (described later) was supported on Android 4.0, which was released around the same time, HLS subsequently became the most common format for delivery.

2.3 The Emergence of H.265 and 4K

The digital processing of video entails a constant battle with data volumes. To illustrate it simply, the higher the image quality, the larger that data will be. The only way to keep data volumes in check is to reduce the picture field (screen size), or make do with lower image quality. The technology for reducing data volume while subjectively keeping the degradation of image quality to a minimum is known as a codec. Codec is a portmanteau of “coder-decoder,” and indicates the playback-side decoding of content coded (or encoded) on the transmission side. This technology is widely used in DVDs and terrestrial digital broadcasting. Currently, the most commonly used codec is a standard known as H.264*2.

H.264 is a standard developed jointly by the ITU and ISO. It was completed in 2003, so it already has a history of over ten years. Over this period CPU performance has increased, and manufacturers have made improvements to their algorithms. H.264 is merely a standard, and the programming of the compression portion is implemented independently by each manufacturer. Partly as a result of this, the video produced by recent encoders has much higher image quality than the early encoders that came out. In the past I have spoken with researchers who said that H.264 still has a lot of room for improvement. However, even with improvements to H.264 performance, it was difficult to bring it in line with 4K.

In response to this, the new H.265*3 codec appeared. During the development of H.265, it was given the name High Efficiency Video Coding (HEVC). H.265 and HEVC refer to the same standard. As the name suggests, this standard enables highly efficient compression, and it was designed based on the advances in CPU technology over the past ten years, in anticipation of the increased computational complexity required for video compression. H.265 is said to be twice as efficient as H.264. Double the efficiency basically means that twice the amount of data can be sent using the same amount of bandwidth as H.264, and file sizes are halved even with quality equivalent to H.264.

*2 H.264 : Advanced video coding for generic audiovisual service (<http://www.itu.int/rec/T-REC-H.264/e>).

*3 H.265 : High efficiency video coding (<http://www.itu.int/rec/T-REC-H.265/e>).

Manufacturers that have codex technology are currently pouring all their efforts into improving the performance of H.265. In other words, it's unlikely that these companies will have any part in the further development of the H.264 codec in the future (as the same developers are involved). As a result, it could be said that the H.264 codec is quite mature at this stage.

2.4 4K

4K generates far more data than HD (High Definition), with a quadruple-sized picture field and double the frame rate. This results in a total of eight times as much data. As-is, it would be utterly impossible to deliver video via broadcast or communications.

The highly-efficient compression of data over communication paths that have only limited bandwidth makes it possible to deliver this data to more users. Aside from the communications world, bandwidth limits also affect broadcasting. To allow for as many channels as possible, each channel should ideally take up minimal bandwidth. In Japan the H.265 standard has been officially adopted for 4K broadcasts, and test broadcasts are already underway. H.254 codecs are also being adopted for a series of communications-based 4K video services that are set to appear in the future. In both cases, H.265 is considered essential for 4K content delivery. This demonstrates how H.265 has become an enabler for 4K.

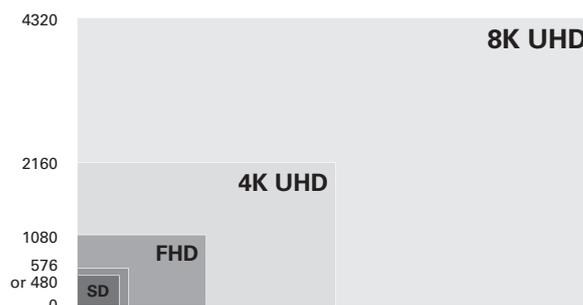


Figure 2: Resolution Comparison

2.5 MPEG-DASH

Meanwhile, Internet-based delivery methods are going through a period of dramatic change. In the 2000s, RTSP (Real Time Streaming Protocol) standardized by the IETF as well as the MMS (Microsoft Media Server) and RTMP (Real Time Messaging Protocol) proprietary methods developed by manufacturers were widely used. These generally utilized UDP to transport data, along with HTTP encapsulation when it was necessary to traverse a firewall. However, in the 2010s there was a shift to delivery methods that use HTTP directly for streaming, or in other words protocols that don't use UDP.

The HTTP Live Streaming (HLS)^{*4} protocol developed by Apple was the pioneering architecture that implemented streaming over HTTP. Apple has submitted this standard to the IETF since 2009, but it has not been adopted by the working group. Apple have documented it as an Internet Draft^{*5} independently, so it has not yet reached RFC status. It uses an extremely simple configuration, in which video data is divided into segments of a few seconds each and placed on a Web server, and then URIs (called manifests) for playing back these segment files are sent to clients in list form. It is not very hard to implement, and many manufacturers have provided support for this method. Meanwhile, HLS remains a standard governed solely by Apple that has not become an RFC Internet standard. There is also no prescribed DRM, which is considered necessary for streaming, and only the most basic encryption has been defined. When completely dependent on it, you have no other option but to follow Apple's policies. In my personal view, Apple sometimes unilaterally takes actions that change the rules. I believe that many companies are looking at whether or not this risk will have an impact on HLS.

For example, HTTP streaming was also adopted in the Smooth Streaming^{*6} protocol proposed by Microsoft. However, since Microsoft began promoting MPEG-DASH, new adoptions of HTTP streaming have slowed down. Smooth Streaming was developed as a part of the Silverlight Rich Internet Application (RIA), but with development of this application terminated, you could say it's reached the end of its life as a format.

*4 HLS: HTTP Live Streaming (<https://developer.apple.com/streaming/>).

*5 HTTP Live Streaming draft-pantos-http-live-streaming (<https://datatracker.ietf.org/doc/draft-pantos-http-live-streaming/>).

*6 IIS Smooth Streaming Technical Overview (<http://www.microsoft.com/en-us/download/details.aspx?id=17678>).

The same applies to the HTTP Dynamic Streaming (HDS)^{*7} proposed by Adobe. If technology developed by a certain company without going through the process of standardization is adopted across the board, there is a risk of vendor lock-in. Furthermore, the streaming industry has only just begun migrating from proprietary protocols to HTTP in the past few years. The desire for a clean technology that allows you to oversee everything yourself may be stronger in the end.

MPEG-DASH^{*8} came about as an antithesis to this kind of situation. Although the technologies are similar, HLS, Smooth Streaming, and HDS each use different segmenting methods and manifests. This means that despite the similarities between these HTTP streaming formats, playback clients and STBs are forced to either implement all of them, or make a choice on which to implement. In light of this situation, the MPEG (Moving Picture Experts Group) ISO/IEC working group decided to develop a standard independent of the formats created by specific manufacturers. A Request For Proposal was issued in 2009, and ultimately the method proposed by the 3GPP (Third Generation Partnership Project) was adopted. MPEG-DASH is short for Dynamic Adaptive Streaming over HTTP. The HTTP standard was adopted for the streaming protocol. It also incorporates adaptive bitrate streaming, which enables streams to be selected based on the client's downstream bandwidth status. By simply preparing streams for multiple bandwidths on the server side, it is possible for clients to switch streams on their own initiative. The adaptive bitrate streaming technique has gained broad acceptance, because there is no need for server-side implementation, and it is possible to respond to subtle changes in the overall client-side environment. MPEG-DASH also meets a range of current content delivery requirements, such as enabling a number DRM methods to be applied.

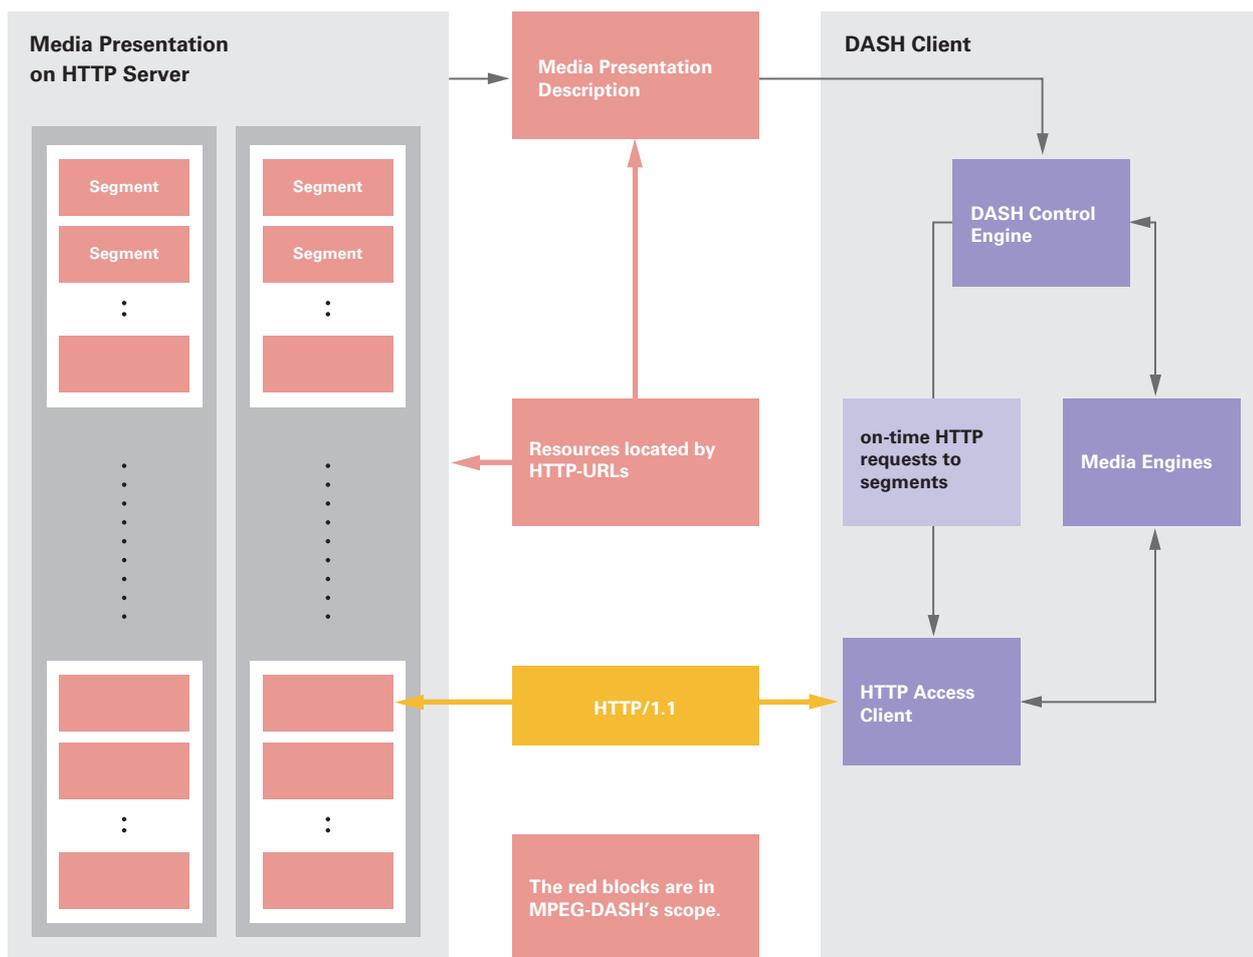


Figure 3: MPEG-DASH

*7 HDS: HTTP Dynamic Streaming (<http://www.adobe.com/products/hds-dynamic-streaming.html>).

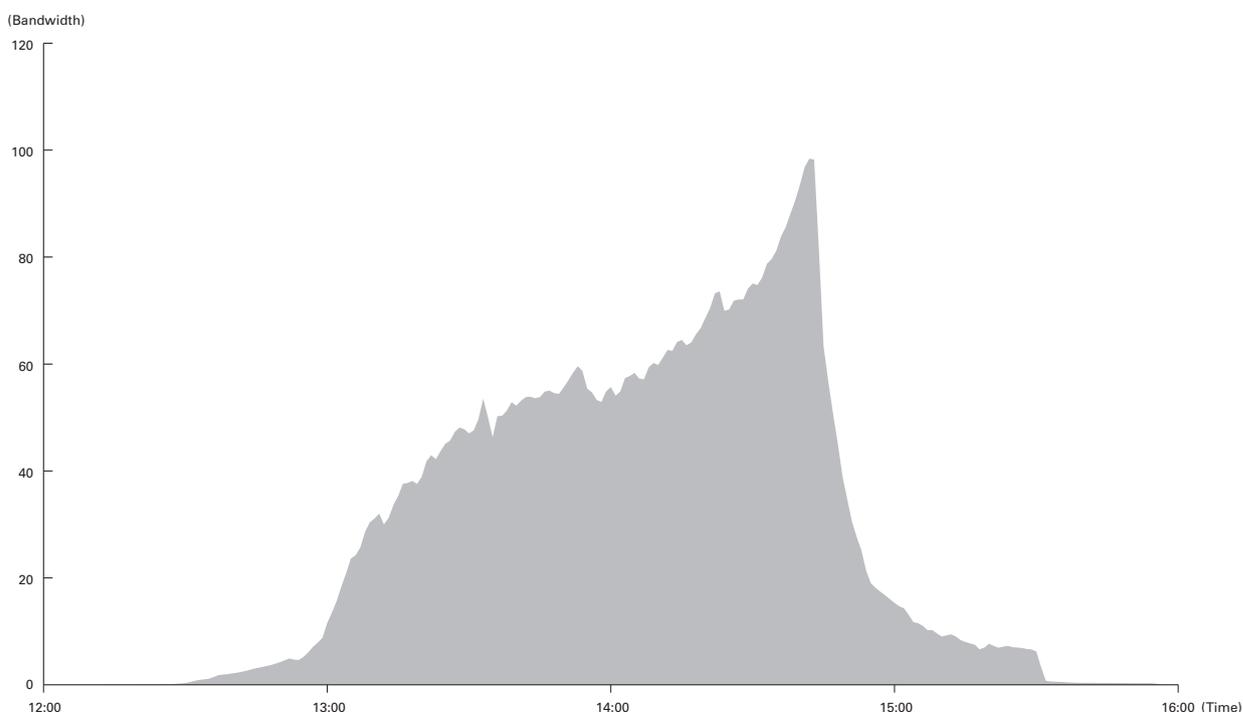
*8 MPEG-DASH "ISO/IEC 23009-1:2014 Dynamic adaptive streaming over HTTP (DASH) -- Part 1: Media presentation description and segment formats": (http://www.iso.org/iso/home/store/catalogue_tc/catalogue_detail.htm?csnumber=65274).

MPEG-DASH caught on in Europe first, and it is now being adopted in rapid succession worldwide. It is currently in the limelight, with its adoption being promoted by an organization called the DASH Industry Forum*⁹, while seminars have also been held for it at large broadcast industry conferences such as the NAB Show and IBC. In Japan, IJ was the first to implement streaming delivery using MPEG-DASH and H.265 at the Spring Festival in Tokyo, which took place in March 2014. This cutting edge trial applied the latest technologies in the field ahead of anyone else, and was achieved by combining technology from a number of manufacturers. The number of encoders and servers that support MPEG-DASH is also on the rise. Going with a standard specification led to a large number of manufacturers adopting the technology, which creates more options for users and promotes competition. It is fair to say that MPEG-DASH is creating a market, and it is on the verge of widespread adoption.

2.6 Summer Koshien Case Study

IJ provides streaming delivery services for the National High School Baseball Championship at Koshien Stadium, which is produced by Asahi Broadcasting Corporation. Every year the special website set up by Asahi Broadcasting Corporation gets large numbers of hits, with most involving the receipt of live streamed content. IJ offers its large-scale CDN services (IJ Contents Delivery Service / Multi Device Streaming / Live) to provide streaming without congestion.

In 2014, we used the HLS (for mobile devices) and HDS (for PCs) delivery protocols. These two protocols were distributed from a single platform. Asahi Broadcasting Corporation encodes the video, then IJ uploads it to the ingest servers. At this stage the RTMP protocol is used. IJ's servers read requests from clients, and switch between HLS and HDS on the fly. The customer doesn't need to prepare two separate streams using the HLS and HDS protocols, and they can continue using their existing RTMP-compatible encoder, so this is a popular feature.



You can see that traffic increases as the innings go by. Peak instantaneous bandwidth on the day was recorded at 108 Gbps.

Figure 4: August 25, 2014 - Traffic Graph for High School Baseball Summer Championship Final Broadcast

*9 DASHIF: DASH Industry Forum (<http://dashif.org/>).

This year the ad insertion technology attracted the most attention. Asahi Broadcasting Corporation triggered the insertion of advertisement videos in between innings. The same ads were inserted for all users this time, but by pairing this technology with the client's IP address, cookies, or technology such as user authentication, it would be easy to personalize the ads shown. As indicated above, delivery to mobile devices will be a large market for streaming business in the future. Mobile devices are extremely personal things, so we believe being able to customize the insertion of ads for devices such as these offers vast business potential.

This was also Asahi Broadcasting Corporation's first attempt at broadcasting the National High School Baseball Championship to mobile devices. Stable delivery is what supports these kinds of attempts. When delivery is disrupted straight away, users spend less time watching a stream. This leads directly to lost business opportunities, so service providers like us must take extreme care and do our utmost to provide stable server operation.

2.7 Conclusion

As shown here, streaming technology is established in standard specifications, and will continue to grow in scale. The advantages of HTTP include improved CDN compatibility, and the fact that it is NAT and firewall friendly. It is also thought to have high affinity with cache servers. Additionally, there are currently moves to evaluate streaming using HTTP/2. Meanwhile, MPEG-DASH has the benefit of facilitating better support from industries that support international standards (through tools or solutions being provided). The same goes for H.265.

IJ will continue striving to stay abreast of the latest technologies such as these and roll out services and solutions. We will also work on collaborations with the IJ Research Laboratory to study technology for larger-scale streaming delivery and more stability.

Author:



Bunji Yamamoto

Mr. Yamamoto is a Senior Engineer in the Product Promotion Department Corporate Planning Section of the Product Division, IJ. He joined IJ Media Communications in 1995. He has worked at IJ since 2005. He is mainly involved with the development of streaming technology. Among his contributions to development of the market is the organization of the Streams-JP Mailing List, which discusses this technology.

The Lagopus SDN Software Switch

Here we explain the capabilities of the new Lagopus software switch in detail, starting with the basics of SDN and OpenFlow.

3.1 SDN and OpenFlow

Those engaged in network-related work are likely to have heard the word SDN come up. Conventional network construction involved manually setting each dedicated hardware device, such as routers and switches, that is installed as a network component. In contrast, the basic concept of Software-Defined Networking (SDN) is to configure settings for network components via software, using information centrally managed through a system called an orchestrator.

According to the Open Networking Foundation (ONF), which is widely promoting SDN and OpenFlow, SDN components can be categorized into three layers (Figure 1).

- **The Application Layer**
- **The Control Layer**
- **The Infrastructure Layer**

The application layer governs the centralized management and user interfaces that serve as the core of SDN. The components of this layer can handle a wide range of management tasks. They cover cloud solutions, or more specifically virtual machines, in addition to SDN. In terms of software this corresponds to orchestrators such as OpenStack and CloudStack.

The control layer serves as a bridge between the application layer and infrastructure layer. From the perspective of the application layer this layer provides network services. Its main role is replacing process requests from the application layer with specific configuration instructions to configure the settings for each component in the infrastructure layer. In terms of software, frameworks such as Floodlight and Ryu are used.

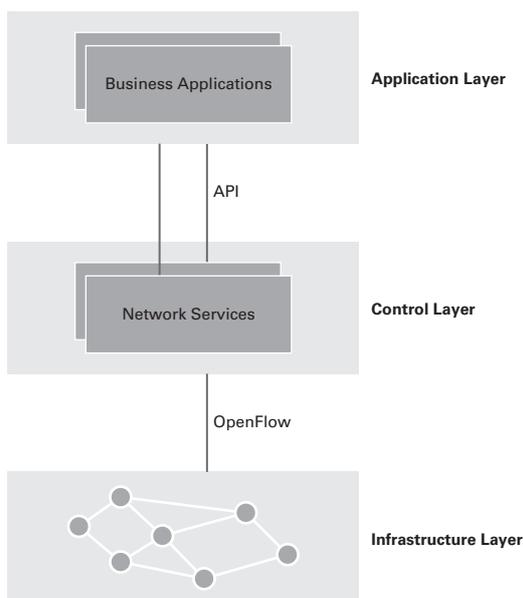


Figure 1: SDN Components

The infrastructure layer, as the name suggests, governs network infrastructure. This corresponds to the components that send and receive customer network packets. In a broad sense, the components in this layer are like routers or switches that enable software control from the control layer. However, in a more limited sense, they are components that can be controlled via vendor-neutral APIs and protocols.

The OpenFlow Switch Specification is a well-known switch specification that enables control via vendor-neutral protocols. A control protocol called the OpenFlow protocol is also defined under this specification for use between controllers and switches. OpenFlow is defined as an open standard, so it is possible to implement hardware or software freely referring to this specification. In addition to hardware switch products, the infrastructure layer components that support OpenFlow include a number of software switches implemented as software.

3.2 OpenFlow

OpenFlow was designed at Stanford University in 2008, and the specification was subsequently expanded and modified, leading to the current open technical specification defined by the ONF. The specification has version numbers assigned, with OpenFlow 1.0 currently the most commonly used. It has been expanded a number of times since then, but although versions 1.1 and 1.2 were developed, upgrades were released before these became popular, so they are not often used. Implementations of OpenFlow 1.3 are often seen, as it is slated for long-term support. At the time of writing the latest version is OpenFlow 1.4, but because the increased complexity of this specification has made it harder to implement, there are currently hardly any implementations that support it. The OpenFlow specification does not take compatibility between versions into consideration, so the versions used between controllers and switches must match.

The concept of OpenFlow begins with the separation of the control plane and the data plane. The data plane processes and transfers packets at high speed according to the rules applied, while the control plane applies these rules. In this case, rules are what distinguish the nature of incoming packets. For OpenFlow, each rule and the corresponding process are paired together in what is known as a flow entry. For example, it is possible to apply a rule for packets sent to TCP port 80 at the destination IP address 192.168.0.1 received over LAN port 1. A collection of flow entries is called a flow table. When a flow entry corresponding to packets processed via OpenFlow is found on a flow table, the associated action is executed. These actions include simply sending the packets to a specified LAN port, as well as modifying the packet headers by altering the destination IP address and inserting a VLAN tag. Actions can be specified freely within the scope of the specification. It is also possible to prepare multiple flow tables, and specify an action that causes the next flow table to be processed.

In addition, you can specify actions that send packets to the control plane, or in other words the controller. The controller receives packets from the switch, and enables more complicated processing to be executed. Additionally, the controller can instruct the switch to send arbitrary packets. Figure 2 gives an overview of these points.

This demonstrates how OpenFlow was designed to enable complicated packet processing while maintaining high speed through coordination between the controller and switch. However, the extension of the OpenFlow specification itself is causing the data plane to become increasingly complicated.

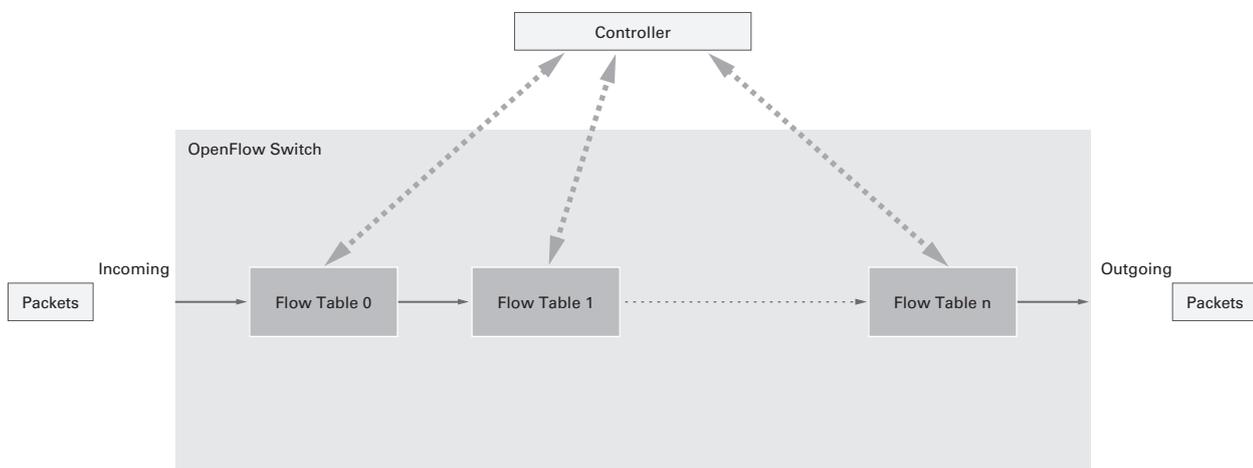


Figure 2: The Concept of OpenFlow

3.3 Software Switches

Broadly speaking, methods for implementing Ethernet switches that support OpenFlow can be grouped into three categories.

The first category is hardware switch products that use ASIC. While these products have large switching capacity, and can provide the performance to support large-scale networks, the number of flow entries and tables that can be handled is considerably limited, and the devices are also expensive. Additionally, although the OpenFlow specification includes items that are required as well as those that can be omitted, it is necessary to note that on hardware switch products the difficult processing of implementations is often skipped.

The second category is products that use network processor units (NPU). These use processors that incorporate special functions for network processing, and are equipped with firmware that takes advantage of these functions. Compared with switches that use ASIC, there are fewer restrictions on the number of flow entries and tables, as well as the items implemented. Also, although their speed is not always on the same level as hardware switches, they can be applied to large-scale networks. They are also expensive, however, and it is difficult to justify the cost for use in medium or small-scale networks.

The third category is software switches. In this case the features of OpenFlow are implemented using programs run on a universal server OS (for example Linux). These can also be run on a virtual machine, meaning they have good compatibility with cloud solutions. While they can provide all the functionality of OpenFlow depending on the implementation, their performance hasn't even been in the same league as ASIC and NPU in the past, and this was a major issue.

The speed when processing 64 byte packets at wire-rate on a 10GbE network interface is 14.88 M frames/sec (fps), which comes to 67 nanoseconds of time allowable per frame. If the processor clock is 2 GHz, each clock cycle is 0.5 nanoseconds. Roughly speaking, if we assume one command per clock cycle, the processing of a packet must be completed within 134 commands or packet transfer at wire-rate will not be possible. The software switch implementations that existed in 2012 were only able to process a few hundred thousand to a million frames per second at most, so the development of a new software switch implementation to resolve performance bottlenecks was investigated. This led to the development of the Lagopus SDN software switch.

3.4 The Lagopus SDN Software Switch

Lagopus is a software switch based on the OpenFlow 1.3 specification. Development began in 2013, and in July 2014 the source code was released under the Apache License. As of October 2014 the latest version is 0.1.1, and development is still ongoing. The main features of Lagopus are listed below.

- **A software switch for x86_64 Linux**
- **Single process, multi-threaded operation (two or more cores required)**
- **Mainly uses the C development language**
- **Based on OpenFlow 1.3.4**
- **Implements most of the options described in the OpenFlow specifications**
- **10 GbE wire-rate performance via packet forwarding with MPLS-VLAN mapping**
- **Supports large-scale configurations with up to 1 million flow entries and 254 tables**
- **High-speed packet processing via the user space using Intel DPDK**
- **Open source software usable for both commercial and non-commercial purposes, with the source code released under the Apache License**

In this report, we examine the performance of this implementation in detail. First, we will describe the main points that have a significant effect on software switch performance.

3.5 Software Switch Bottlenecks

3.5.1 Receive Interrupt and Transmission Complete Interrupt Processing

When the network interface card (NIC) detects the receipt of packets while the CPU is executing various processes, the NIC stores the packet byte string in buffer memory and issues an interrupt to the CPU. When the CPU detects an interrupt from the NIC, it temporarily saves the current process state (context) before dealing with the interrupt. The minimum amount of interrupt processing is executed, such as retrieving the packet information stored in the buffer memory, and putting it on the receive queue of the network stack. When this processing is complete, the saved context is restored, and processing from before the interrupt was issued is resumed. The act of switching from one context to another is called a context switch, and this creates a large load when transferring packets at high speed. Because generating an interrupt for each packet has a significant effect on performance, there are NIC functions designed to reduce the load by issuing interrupts once a certain number have accumulated.

3.5.2 Copying Packets to Memory

Received packets are first stored to the memory managed in the OS kernel space, but this is not normally directly readable by programs in the user space. The read (2) or recv (2) system calls copy the content of received packets to a memory space secured in advance by the user program. Conversely, the write (2) or send (2) system calls copy the content prepared by the user program to the memory space in the kernel. Because 14.88 million frames are sent or received each second at a wire-rate of 10GbE, for 64 byte packets a total of 1.9 GB (14.88 million x 64 x 2) of memory copying occurs per second. DDR3-1600 memory has a theoretical transfer rate of 12.8 GB per second, so even simply applying these results in approximately 15% of the total time being consumed for memory copying alone.

3.5.3 OS Process Scheduling

A number of processes are run in the user space of an OS, but as you can see by looking at the ps commands, even with the current popularity of multi-core processors we are not yet at the point where enough CPU cores are available. Consequently, the OS switches between running processes as appropriate to make it appear as if they are all running at once. Because switching is carried out at the millisecond level, it takes place too quickly for users to notice during normal use. However, if this kind of OS process scheduling occurs while a software switch is running, performance is naturally greatly impacted.

3.5.4 Memory Block Allocation and De-allocation, and Page Faults

Regarding the memory blocks used when reading or writing, it is not efficient to allocate them each time packets are received and de-allocate them each time transmission is completed. Considering that the next packets will be received again straight away, reusing the same space can alleviate performance degradation. However, not all memory in the user space is placed in actual memory. The OS maps user space memory and actual memory as necessary. This happens the instant a program attempts to access that memory. This system is called on demand paging, and the process triggers an exception the moment that a program attempts to access memory. Memory mapping is then executed during exception handling, and the program continues processing as if nothing has happened after it is restored. This exception is known as a page fault. It is obvious there is a large difference in processing time between cases in which memory mapping is executed when a page fault occurs, and cases in which memory mapping has already been carried out.

3.5.5 Exclusive Access to Resources

Above we explained that packet processing is executed by referring to a flow table under OpenFlow, but naturally it is also possible to add, change, and delete entries to a flow table during packet processing. However, a flow entry currently being referenced that matches a received packet must never be deleted. For this reason, the execution of packet processing and the addition, change or deletion of entries to a flow table must be mutually exclusive. In addition to flow tables, other resources that require exclusive access include LAN port counter information. A mechanism called a readers-writer lock is usually used for this, but this lock generates comparatively heavy processing, and repeated locking and unlocking of each packet would have a significant effect on transfer performance.

3.6 Resolving Bottlenecks With Intel DPDK

There are already a number of attempts to resolve several of the bottlenecks mentioned above, but here we will discuss the Intel DPDK initiative used with Lagopus. Intel DPDK (Data Plane Development Kit) is a library that provides useful API for programming data planes. It serves as a framework for enabling high-speed packet I/O. At the time of writing, the latest version is 1.7.1. The main features of Intel DPDK are as follows.

- **Runs on x86 and x86_64 Linux as well as FreeBSD**
- **Provided under a BSD License, so it can be used for both commercial and non-commercial purposes**
- **Poll Mode Driver (PMD) enables sending and receipt of packets without using interrupts**
- **Zero-copy via user space I/O**
- **OS scheduling avoidance designed for multi-core based on core fixing**
- **Page fault deletion using hugepage**
- **Optimized libraries including Longest Prefix Match, hash, and ring buffer**
- **Many sample applications included**

3.7 Performance Improvements in Lagopus

The use of Intel DPDK with Lagopus resolves many performance bottlenecks. In particular, PMD and core fixing contribute greatly to improved performance. The thread configuration of Lagopus is also quite distinctive (Figure 3).

The Lagopus data plane is divided into threads for handling packet I/O and threads for handling OpenFlow processing. Furthermore, it is possible to specify incoming (RX) and outgoing (TX) threads for I/O, and multiple threads for OpenFlow processing (workers) upon startup. These threads are each fixed to run on separate cores to prevent other processes being allocated to them, maintaining performance. A ring buffer provided by DPDK that does not require lock processing is used for communication between threads.

Bottlenecks are mainly resolved using Intel DPDK for basic packet I/O, but for OpenFlow processing Lagopus has its own performance improvements. The internal structure of the Lagopus flow table is shown in Figure 4.

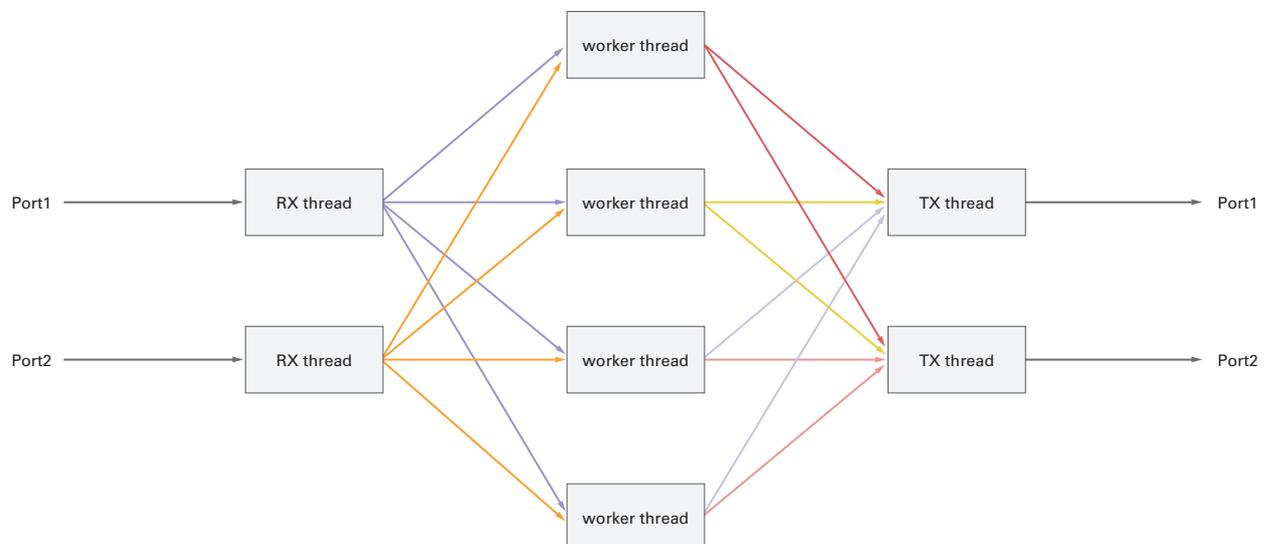


Figure 3: Lagopus Thread Configuration

Because it is necessary to output flow table details in response to requests from the controller, a flat table in line with the OpenFlow specification is available, but Lagopus also provides another table with a structure suitable for flow lookup. This structure is currently fixed, and not every flow table enables optimal lookup. However, it facilitates high-speed lookup of typical packet headers, such as branching by VLAN ID, MPLS label, or IPv4 destination address. A flow cache system is also implemented to cache the information for packets that have already been looked up once. When the same packets are received, flow entry information is obtained from the cache and the lookup itself is skipped.

Due to synergy from the above performance improvements, Lagopus was able to achieve 10GbE wire-rate performance (14.88 Mfps) for packet forwarding with packet header modification via OpenFlow, in a test environment using Intel Xeon processors with a flow table that contained 100,000 entries. The current focus for Lagopus is improving performance with more complex flow entries, maintaining performance when processing more ports simultaneously, and providing support for higher speed interfaces.

3.8 Conclusion

The Lagopus SDN software switch is highly compliant with the OpenFlow specification, and it enables dramatically improved performance in comparison to conventional software switches. Lagopus is currently prototype quality software that has not been sufficiently tested or validated outside a test environment. However, development continues with the assumption that it will see widespread adoption in real environments. If you are interested in learning more, please take a look at the website (<http://lagopus.github.io>).

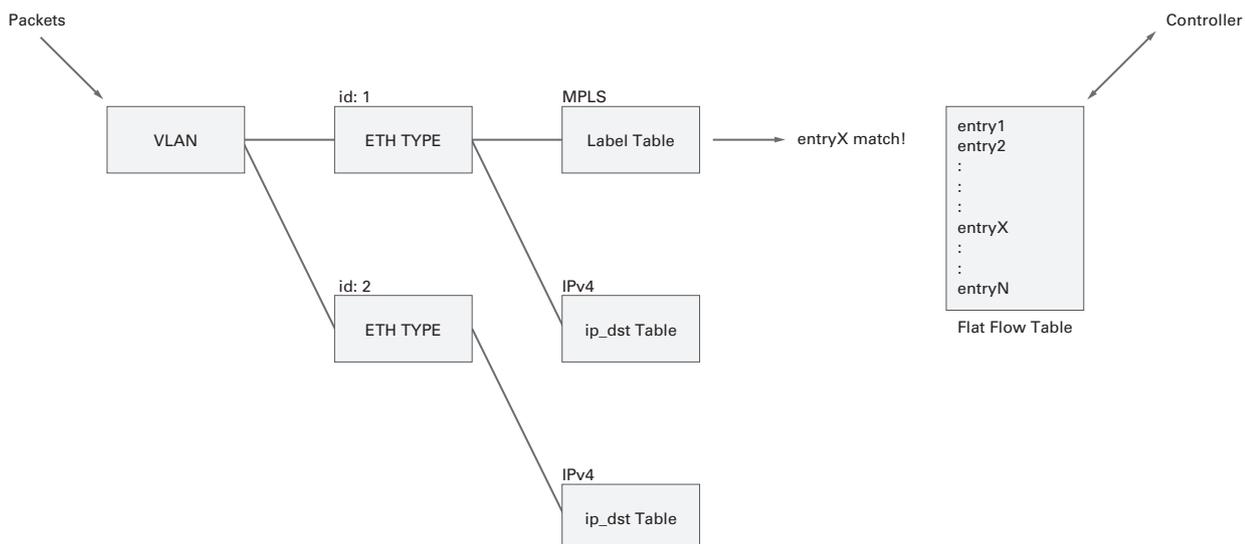


Figure 4: The Internal Structure of Lagopus Flow Tables

Author:



Masaru Oki
Research & Development Department, Stratosphere Inc.

About Internet Initiative Japan Inc. (IIJ)

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

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

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

Internet Initiative Japan Inc.

Address: Iidabashi Grand Bloom, 2-10-2 Fujimi, Chiyoda-ku, Tokyo 102-0071, Japan
Email: info@ij.ad.jp URL: <http://www.ij.ad.jp/en/>

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.

©2008-2014 Internet Initiative Japan Inc. All rights reserved.

IIJ-MKTG020WA-1412CP-00001PR