

Internet Infrastructure Review Vol.26

February
2015

Infrastructure Security

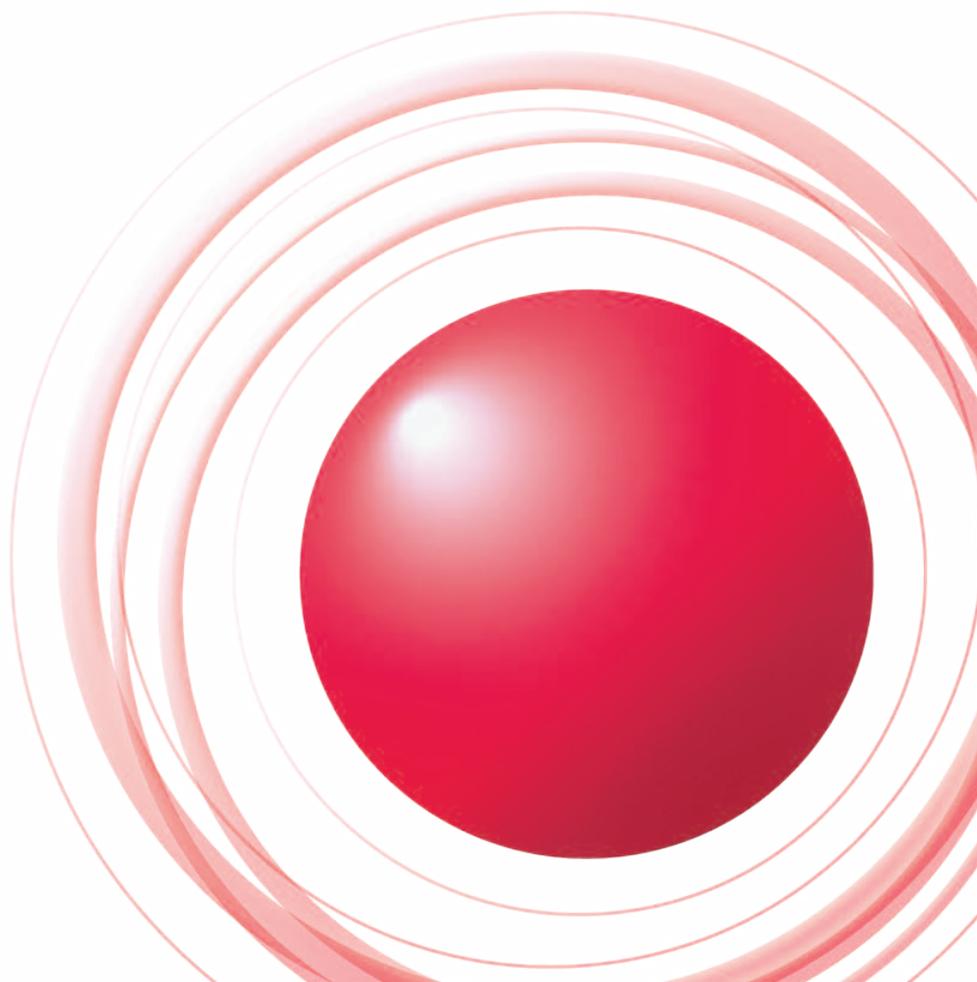
Countermeasures Against the Alteration of Domain Name Registration Information

Cloud Computing Technology

Introducing the IIJ GIO Cloud Service - Part 2

Storage Technology

UKAI Virtual Disk Storage



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 Countermeasures Against the Alteration of Domain Name Registration Information	18
1.4.2 The openioc_scan Plug-in That Scans for Threats Lurking in Device Memory	21
1.4.3 ID Management Technology	24
1.5 Conclusion	27
2. Cloud Computing Technology	28
2.1 Introduction	28
2.1.1 Cloud Market Trends	28
2.1.2 Service Overview	28
2.2 Large-Scale Infrastructure Technology Initiatives	29
2.2.1 Infrastructure Development	29
2.2.2 Infrastructure Operation	32
2.3 Conclusion	33
3. Storage Technology	34
3.1 Infrastructure with Virtualization as a Prerequisite	34
3.2 The Necessity of Disk Storage for Virtual Machines	35
3.3 UKAI: Location Aware Virtual Disk Storage	36
3.4 Implementing UKAI	37
3.5 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

In recent years, a “cloud-first” mindset has been instilled in companies, with the use of cloud services considered above all else when building or renewing corporate IT systems. According to a report called “Japan Storage in the Cloud 2014–2018 Forecast and 2013 Analysis” published by IDC in October of last year, the market scale of pay-as-you-go storage services offered as public cloud services (storage in the cloud), which was 15,819,000,000 yen in 2013, is forecast to subsequently expand at an average annual growth rate of 26.7%, reaching 51,560,000,000 yen in 2018.

Looking at the status of facilities for IJ’s IJ GIO cloud service, the physical storage capacity required for services increased by approximately 40% between the end of 2013 and the end of 2014, and is starting to demonstrate more growth than server numbers. Even for services with surging demand such as this, it is necessary to maintain overall system availability for the infrastructure providing them, while also dealing with continual expansion and renewal of facilities, handling urgent failures and faults, and implementing vulnerability countermeasures reliably. To do this, at the facility design stage it is imperative to thoroughly evaluate matters ahead of time, including which technology and products to adopt, and which architecture to build, to secure features such as fault-tolerance, maintainability, and scalability.

This report discusses the results of the various ongoing surveys and analysis activities that IJ, as a service provider, carries out to support the Internet and cloud infrastructure, and enable our customers to continue to use them 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 October 1 to December 31, 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, discussing countermeasures against the alteration of domain name registry registration information, and looking at a tool that scans for traces (IOCs) of threats lurking in the memory of a device. In addition, we provide an explanation of ID management technology.

In the “Cloud Computing Technology” section, we give an overview of the new IJ GIO cloud service that IJ provides, and summarize the deployment status and usage trends for the service facilities that support it. We also examine ongoing initiatives for the stable operation of service infrastructure, such as approaches to the construction of large-scale infrastructure, the renewal of facilities to maintain infrastructure, and processes such as failure and fault handling.

In the “Storage Technology” section, we present our research into the “UKAI” storage system that enables flexible control over the actual data location of virtual disks, for which location management is difficult as service infrastructure moves toward virtualization and the cloud. We first give an overview of the status of virtualization infrastructure that UKAI is designed around, then discuss issues that need to be overcome, architecture for resolving these issues, and details of a prototype implementation.

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.

Countermeasures Against the Alteration of Domain Name Registration Information

In this report, we discuss countermeasures against the alteration of domain name registration information, and examine the openioc_scan plug-in that scans for threats lurking in the memory of a computer. We also take a look at ID management technology.

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 October 1 through December 31, 2014. In this period a number of hacktivism-based attacks were once again carried out by Anonymous, and a series of targeted attacks on companies such as media outlets were discovered. There were also incidents of suspicious messages being displayed as well as malware infections originating from alterations and domain hijackings caused by attacks on domain name registries, with companies in Japan also being affected. In November, there was an incident in which a movie-related company in the United States had its corporate IT system rendered inoperable. A large quantity of information was also stolen, and some of this was released to the public. 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 October 1 and December 31, 2014. Figure 1 shows the distribution of incidents handled during this period*1.

■ 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. In October there were a number of defacements and information leaks affecting websites of the Chinese government and the Hong Kong executive branch in relation to demonstrations in Hong Kong (OpHongKong). In November, the websites of a number of government institutions in the Philippines were defaced in protest against the government. In Turkey, the website of a power transmission company was compromised by a party protesting against the Turkish government. There

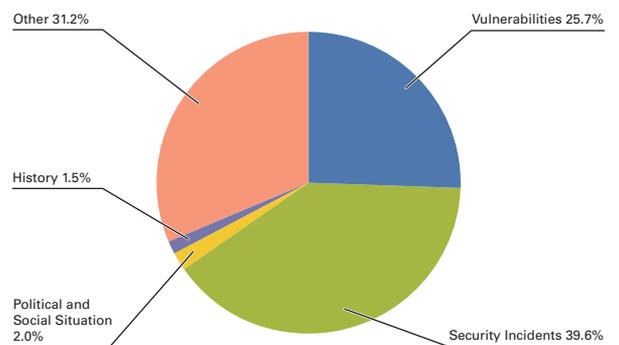


Figure 1: Incident Ratio by Category (October 1 to December 31, 2014)

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

were also a number of website defacements and compromises in Turkey thought to be the work of another party. Attacks continue to be made on a number of sites for the Israeli government and private-sector businesses in relation to the conflict in Palestinian-controlled Gaza Strip. Similarly, attacks related to conflict were also made between countries such as Russia and Ukraine, Pakistan and Indonesia, and India and Pakistan. Other attacks by hacktivists such as Anonymous continued on government and government-related websites around the world. There were also ongoing hijacking incidents targeting the SNS accounts of government institutions and companies such as media outlets.

■ Vulnerabilities and their Handling

During this period fixes were released for Microsoft's Windows^{*2*3*4*5}, Internet Explorer^{*6*7*8}, and Microsoft Office^{*9}. In December, a lot of interest was generated when Google disclosed a vulnerability^{*10} it had discovered in Windows 8.1 Update with no fix available. This was due to Google's policy of automatically disclosing vulnerability information 90 days after it is reported in some cases, but some expressed concern at Google making this information public^{*11}. Updates were also made to Adobe Systems' Flash Player, Reader, and Acrobat. A vulnerability in JustSystems Corporation's Ichitaro that could allow execution of arbitrary program code by a third party was discovered and fixed. A quarterly update was provided for Oracle's Java SE, fixing many vulnerabilities. Several of these vulnerabilities were exploited in the wild 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. An SQL injection vulnerability was also discovered and fixed in the Drupal CMS. A number of vulnerabilities, including an XSS vulnerability, were fixed in WordPress as well. Cases in which these vulnerabilities were exploited to actually steal administrative privileges have been confirmed^{*12}. A vulnerability that could allow an external party to cause abnormal operations on a server or stop a service has been fixed in several pieces of DNS software, such as BIND and Unbound. A vulnerability that could allow arbitrary code execution with ntpd execution privileges using specially-crafted packets was discovered and fixed in the ntpd program used for time synchronization.

In October, a new method for attacking the SSLv3 protocol used with SSL/TLS servers called the POODLE attack (Padding Oracle On Downgraded Legacy Encryption attack) was disclosed^{*13}. Because this vulnerability was caused by issues in the specifications of the protocol itself, the measures taken included fixes to disable SSLv3 in major Web browsers and servers, and the release of information on how to configure settings to prevent SSLv3 communications.

■ Attacks on Domain Name Registries

During this period, there were a number of incidents in which registration information at domain name registries and registrars was illegally altered in attacks, resulting in the redirection of users to malicious sites. In October, the PANDI ccTLD registry that manages Indonesia's .id domains was compromised, and visitors to Google's local site were redirected to another site. On a major video-sharing website it was confirmed that fake ads were used in attacks that attempted to infect

*2 "Microsoft Security Bulletin MS14-058 - Critical: Vulnerabilities in Kernel-Mode Driver Could Allow Remote Code Execution (3000061)" (<https://technet.microsoft.com/library/security/ms14-058>).

*3 "Microsoft Security Bulletin MS14-060 - Important: Vulnerability in Windows OLE Could Allow Remote Code Execution (3000869)" (<https://technet.microsoft.com/library/security/ms14-060>).

*4 "Microsoft Security Bulletin MS14-064 - Critical: Vulnerabilities in Windows OLE Could Allow Remote Code Execution (3011443)" (<https://technet.microsoft.com/library/security/ms14-064>).

*5 "Microsoft Security Bulletin MS14-066 - Critical: Vulnerability in Schannel Could Allow Remote Code Execution (2992611)" (<https://technet.microsoft.com/library/security/ms14-066>).

*6 "Microsoft Security Bulletin MS14-056 - Critical: Cumulative Security Update for Internet Explorer (2987107)" (<https://technet.microsoft.com/library/security/ms14-056>).

*7 "Microsoft Security Bulletin MS14-065 - Critical: Cumulative Security Update for Internet Explorer (3003057)" (<https://technet.microsoft.com/library/security/ms14-065>).

*8 "Microsoft Security Bulletin MS14-080 - Critical: Cumulative Security Update for Internet Explorer (3008923)" (<https://technet.microsoft.com/library/security/ms14-080>).

*9 "Microsoft Security Bulletin MS14-081 - Critical: Vulnerabilities in Microsoft Word and Microsoft Office Web Apps Could Allow Remote Code Execution (3017301)" (<https://technet.microsoft.com/library/security/ms14-081>).

*10 This vulnerability was fixed in January 2015 with "Microsoft Security Bulletin MS15-008 - Important: Vulnerability in Windows Kernel-Mode Driver Could Allow Elevation of Privilege (3019215)" (<https://technet.microsoft.com/library/security/ms15-008>).

*11 For example, problems were pointed out in the following SOPHOS nakedsecurity blog post. "Zero-day in Windows 8.1 disclosed by Google" (<https://nakedsecurity.sophos.com/2015/01/03/zero-day-in-windows-8-1-disclosed-by-google/>).

*12 For example, with Drupal an alert was issued due to the fact that attacks had begun a few hours after the fix was released on October 16. "Drupal Core - Highly Critical - Public Service announcement PSA-2014-003" (<https://www.drupal.org/PSA-2014-003>).

*13 See "1.4.2 The POODLE Attack" in Vol.25 of this report (http://www.ij.ad.jp/en/company/development/iir/pdf/iir_vol25_EN.pdf) for more information.

October Incidents

1	O 4th: The Federal Communications Commission (FCC) ordered a major hotel chain to pay a \$600,000 fine for intentionally blocking the Wi-Fi communications of customers and promoting their paid Wi-Fi system. "Marriott to Pay \$600K to Resolve WiFi-Blocking Investigation" (http://www.fcc.gov/document/marriott-pay-600k-resolve-wifi-blocking-investigation-0). Regarding this issue, a warning was issued in January once again calling for business operators to not intentionally block tethering using Wi-Fi routers or smartphones. "WARNING: Wi-Fi Blocking is Prohibited" (http://www.fcc.gov/document/warning-wi-fi-blocking-prohibited).
2	
3	
4	S 5th: An incident occurred in which the PANDI ccTLD registry for Indonesia's .id domains was accessed by an unknown party without authorization, and Google's site was hijacked.
5	S 10th: U.S. security company Volexity reported that cyber attacks were occurring in relation to demonstrations in Hong Kong. Among these, it was identified that a number of sites in Japan, including the Web server of a Japanese newspaper publisher, had been altered to act as malicious sites or redirect users to malicious sites. See the following blog post for more information: "Democracy in Hong Kong Under Attack" (http://www.volexity.com/blog/?p=33).
6	
7	S 10th: It came to light that data such as customer credit card information had leaked from a number of U.S. retailers due to malware. For example, it is reported that the store payment system at major discount chain Kmart had been infected with malware. Sears Holdings Corporation "Kmart Investigating Payment System Intrusion" (http://searsholdings.mediaroom.com/index.php?s=16310&item=137317).
8	
9	O 10th: The JPCERT Coordination Center issued an alert due to an increase in scans to TCP port 10000, which is used by the Webmin Web-based system administration tool. "JPCERT/CC Alert 2014-10-10 Alert regarding increase in scans to TCP port 10000" (http://www.jpCERT.or.jp/english/at/2014/at140038.html).
10	
11	V 15th: Microsoft published their Security Bulletin Summary for October 2014, and released a total of eight updates, including three critical updates such as MS14-056 and MS14-058, as well as five important updates such as MS14-060. "Microsoft Security Bulletin Summary for October 2014" (https://technet.microsoft.com/library/security/ms14-oct).
12	V 15th: Oracle released their quarterly scheduled update for a number of products including Oracle, fixing a total of 154 vulnerabilities, including 25 in Java SE. "Oracle Critical Patch Update Advisory - October 2014" (http://www.oracle.com/technetwork/topics/security/cpuoct2014-1972960.html).
13	
14	V 15th: Google disclosed the new POODLE (Padding Oracle On Downgraded Legacy Encryption) method for attacking the CBC cryptographic algorithm used by SSL 3.0. See the following article from the Cryptographic protocol Evaluation toward Long-Lived Outstanding Society (CELLOS) for more information: "[2014-10-15] On a new attack on SSLv3 specification called POODLE attack" (https://www.cellos-consortium.org/index.php?PoodleAttack_20141015).
15	
16	V 15th: A number of vulnerabilities in Adobe Flash Player that could allow arbitrary code execution were discovered and fixed. "Security updates available for Adobe Flash Player" (http://helpx.adobe.com/security/products/flash-player/apsb14-22.html).
17	S 15th: Security warnings caused by Google Safe Browsing were displayed on some websites that featured a Hatena Bookmark button. Hatena Bookmark Development Blog, "Regarding security warnings on some sites featuring a Hatena Bookmark button" (http://bookmark.hatenastaff.com/entry/2014/10/18/021046) (in Japanese).
18	
19	V 16th: An SQL injection vulnerability (CVE-2014-3704) was discovered and fixed in the Drupal CMS application. "SA-CORE-2014-005 - Drupal core - SQL injection" (https://www.drupal.org/SA-CORE-2014-005).
20	S 16th: Trend Micro issued an alert due to the confirmation of attacks that redirect users to a malware distribution site from fake ads on YouTube. Trends Micro Security Intelligence Blog, "YouTube Ads Lead To Exploit Kits, Hit US Victims" (http://blog.trendmicro.com/trendlabs-security-intelligence/youtube-ads-lead-to-exploit-kits-hit-us-victims/).
21	
22	O 17th: The Ministry of Economy, Trade and Industry compiled and published its "Guidelines concerning the Online Services for Consumers regarding Notification on the Use of Private Data and Gaining Approval from the Consumer for Data Usage," to serve as reference for companies offering services overseas. "Announcement of guidelines for providing information and explanations in consideration of the privacy of consumers on online services" (http://www.meti.go.jp/press/2014/10/20141017002/20141017002.html) (in Japanese).
23	
24	O 17th: The National Police Agency issued an alert due to an increase in SSDP reflection attacks. "Alert regarding SSDP reflector attacks using UPnP compatible network devices as stepping stones" (http://www.npa.go.jp/cyberpolice/detect/pdf/20141017.pdf) (in Japanese).
25	
26	O 21st: Apple released a statement regarding issues with the Spotlight search function in OS X Yosemite sending user's current location and search information to them. "OS X Yosemite: Spotlight Suggestions" (http://support.apple.com/kb/PH18943?viewlocale=en_US).
27	V 22nd: Microsoft announced there was a vulnerability in Microsoft OLE (CVE-2014-6352) without a fix available that could allow remote code execution. "Microsoft Security Advisory 3010060 Vulnerability in Microsoft OLE Could Allow Remote Code Execution" (https://technet.microsoft.com/library/security/3010060). This vulnerability was fixed in "Microsoft Security Bulletin MS14-064 - Critical: Vulnerabilities in Windows OLE Could Allow Remote Code Execution (3011443)" (https://technet.microsoft.com/library/security/ms14-064).
28	
29	O 31st: The Ministry of Internal Affairs and Communications made an announcement regarding their initiatives to improve the mobile usage environment, such as removing SIM locks, promoting MVNOS, and allocating frequency spectrum for 4G and other high speed communications. The goal of these initiatives is to stimulate the economy and reduce the national burden through the creation of new mobile-based business. "Announcement of Mobile Recreation Plan" (http://www.soumu.go.jp/main_sosiki/joho_tsusin/eng/Releases/Telecommunications/141031_05.html).
30	
31	

[Legend]

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

*Dates are in Japan Standard Time

users with ransomware. In this incident, it was identified that the attackers somehow altered the DNS information of the Polish government's domain and made it appear as a legitimate site. Incidents in which attempts were made to redirect users to malicious sites or infect certain users with malware have also occurred at major newspaper publishers and SNS services in Japan, with the cause thought to be the unauthorized alteration of domain name registration information. A U.S. security firm has reported that these attacks were related to the demonstrations held in Hong Kong. In November there were also incidents in which messages and images from an unknown party calling themselves the Syrian Electronic Army were displayed on a number of major sites in Japan. In these incidents, an SNS coordination service on these sites was used to redirect users to another site thought to have been prepared by the attacker, who had rewritten registration information related to the name server after gaining unauthorized access to the registry. This had an impact on many sites, including major sites in Japan and overseas that were using the corresponding service. Because there had been a number of similar incidents in which Japanese companies were affected, several organizations including the JPCERT Coordination Center issued alerts. See "1.4.1 Countermeasures Against the Alteration of Domain Name Registration Information" for more information.

■ Government Agency Initiatives

Government agency initiatives included passing the "Bill on Cyber Security," which had been carried over from the previous Diet session, at the Lower House plenary session in November. This is designed to strengthen systems for promoting cyber security and functionality to bolster security at government agencies, etc., such as establishing a Cyber Security Strategic Headquarters, creating security strategy plans, formulating security standards for government agencies, and investigating security incidents that occur there. It is also expected to lead to more collaborations between the state and private-sector companies, such as the promotion of security at critical infrastructure providers, and the reinforcement of public-private coordinated countermeasures.

In light of this, the 41st assembly of the Information Security Policy Council was held, and policy for initiatives related to improving the functions of systems for promoting cyber security in Japan was decided on*¹⁴. Other specific initiatives included the establishment of the "Information Security Social Promotion Association" in November. This is aimed at improving the standards of information security in Japan by enhancing the interest, understanding, and response capabilities of the nation as a whole with regard to information security. Also in November, regarding the so-called Japanese version of the NCFTA*¹⁵ discussed at the National Police Agency's General Security Countermeasures Conference, the "Japan Cybercrime Control Center (JC3)" commenced operation. Its participants include private-sector companies, academic research institutions, and investigative organizations. JC3 conducts information sharing and analysis, as well as research and development into countermeasures, in relation to cyberspace threats that participant organizations face. It also provides training, and carries out activities to promote a safe and secure cyberspace through initiatives and collaboration with overseas institutions such as U.S. NCFTA*¹⁶.

■ Website Alteration Techniques

In October, a website operated by a newspaper publisher in Japan was compromised and exploited as a phishing site. Due to this incident, the corresponding site was shut down and an investigation carried out, but service resumed over a month later in December after the investigation was completed and the security of the site was verified. In December, there was an incident of unauthorized access at ICANN (Internet Corporation for Assigned Names and Numbers), which manages and coordinates various Internet resources such as domains and IP addresses. After it was confirmed that a number of systems had been accessed, including the Centralized Zone Data System (czds.icann.org), it was announced that they had taken the step of disabling passwords. In this incident, the information used included passwords that had leaked through a phishing email attack in November in which the attacker posed as an internal staff member. Similarly, unauthorized access was made to the website of Internet Systems Consortium (ISC), which develops the BIND DNS server. This led to alterations that redirected users to a malware distribution site. The alterations made in this incident exploited a vulnerability in WordPress.

*14 National Information Security Center, "41st Assembly (held on a rotating basis) (November 25, 2014)" (<http://www.nisc.go.jp/conference/seisaku/index.html#seisaku41>) (in Japanese).

*15 Because it was specified in the cyber security strategy decided on by the Information Security Policy Council in 2013 that establishment of a Japanese version of NCFTA should be looked into, this matter was examined at the National Police Agency's General Security Countermeasures Conference (<http://www.npa.go.jp/cyber/csmeeting/index.html>) (in Japanese).

*16 NCFTA (National Cyber-Forensics & Training Alliance) (<http://www.ncfta.net/Index.aspx>) is a U.S. nonprofit organization with members comprised of law enforcement agencies such as the FBI, private-sector companies, and academic institutions. It gathers and analyzes information related to cybercrime, and conducts training for the staff of investigative organizations.

November Incidents

1	<p>O 5th: Alerts were issued by a number of organizations including the JPCERT Coordination Center due to multiple domain name hijacking incidents, in which the registration information of .com domain names used by organizations in Japan was illegally altered and name server information prepared by the attacker added.</p> <p>"JPCERT/CC Alert 2014-11-05 Alert on domain name hijacking by altering registration information" (http://www.jpccert.or.jp/english/at/2014/at140044.html).</p>
2	
3	
4	<p>O 6th: The Bill on Cyber Security was passed at the Lower House plenary session.</p> <p>"Bill on Cyber Security" (http://www.shugiin.go.jp/internet/itdb_gian.nsf/html/gian/honbun/houan/g18601035.htm) (in Japanese).</p>
5	
6	<p>S 7th: The FBI announced they had uncovered the Tor-based Silk Road 2.0 crime site that dealt in illicit drugs and arrested the operator.</p> <p>FBI, "Manhattan federal court to the operator of charge Silk Road 2.0 web site" (http://www.fbi.gov/newyork/press-releases/2014/operator-of-silk-road-2.0-website-charged-in-manhattan-federal-court).</p>
7	
8	<p>S 7th: The European Police Office (Europol) announced it had uncovered a number of crime sites operated on Tor in coordination with law enforcement agencies in 16 countries including the United States and European nations.</p> <p>EUROPOL, "GLOBAL ACTION AGAINST DARK MARKETS ON TOR NETWORK" (https://www.europol.europa.eu/content/global-action-against-dark-markets-tor-network).</p>
9	<p>S 11th: The United States Postal Service (USPS) announced it had been accessed by an unknown party without authorization, and the personal information of over 800,000 employees may have leaked.</p> <p>"Postal Service Statement on Cyber Intrusion Incident" (http://about.usps.com/news/fact-sheets/scenario/media-statement-final.pdf).</p>
10	
11	<p>V 12th: Microsoft published their Security Bulletin Summary for November 2014, and released a total of fourteen updates, including four critical updates such as MS14-064 and MS14-065, as well as eight important updates.</p> <p>"Microsoft Security Bulletin Summary for November 2014" (https://technet.microsoft.com/library/security/ms14-nov).</p>
12	
13	<p>V 12th: A number of vulnerabilities in Adobe Flash Player that could allow arbitrary code execution were discovered and fixed.</p> <p>"Security updates available for Adobe Flash Player" (http://helpx.adobe.com/security/products/flash-player/apsb14-24.html).</p>
14	<p>V 13th: A vulnerability in JustSystems Corporation's Ichitaro that could allow arbitrary code executions was discovered and fixed.</p> <p>JVN, "JVN#16318793 Ichitaro series vulnerable to arbitrary code execution Critical" (https://jvn.jp/en/jp/JVN16318793/).</p>
15	
16	<p>O 13th: The JC3: Japan Cybercrime Control Center was launched and began operations. This initiative is aimed at investigating cybercrime and reducing or neutralizing its threat through information sharing between industrial, government, and academic organizations and collaboration with overseas institutions, to deal with threats in cyberspace.</p> <p>"Establishment of 'Japan Cybercrime Control Center,' a New Organization for Fighting Cybercrime" (https://www.jc3.or.jp/media/pdf/pressreleaseEnglish.pdf).</p>
17	
18	<p>O 17th: The inaugural meeting of the Information Security Social Promotion Association was held. The aim of this association is to build a safe and secure society through the construction of an information distribution network by national and regional organizations in industry, government and academia, allowing them to collaborate and coordinate.</p> <p>National Information Security Center (NISC), "Regarding the Inaugural Meeting of the Information Security Social Promotion Association (provisional name)" (http://www.nisc.go.jp/press/pdf/syakaisuishinkyougikai20141113.pdf) (in Japanese).</p>
19	
20	
21	<p>S 19th: The Tokyo Metropolitan and 19 prefectural police departments established a joint investigation headquarters and carried out a simultaneous search due to the illegal operation of proxy servers using unlawfully obtained IDs and passwords. This resulted in the arrest of a number of people suspected of violating the Unauthorized Computer Access Law.</p>
22	
23	<p>V 21st: A number of vulnerabilities that could allow sites to be compromised, including XSS vulnerabilities, were discovered and fixed in the WordPress CMS application.</p> <p>"WordPress 4.0.1 Security Release" (https://wordpress.org/news/2014/11/wordpress-4-0-1/).</p>
24	
25	<p>S 21st: A number of people including minors were charged with violating the Act on Punishment of Activities Relating to Child Prostitution and Child Pornography, and the Protection of Children (display in public), or were referred to a child guidance center, for proliferating child pornography images posted to Twitter by retweeting them.</p>
26	<p>S 24th: An incident occurred in which a movie company in the U.S. was compromised by an unknown party, leading to their entire system being shut down.</p>
27	
28	<p>V 26th: A number of vulnerabilities in Adobe Flash Player that could allow arbitrary code execution were discovered and fixed.</p> <p>"Security updates available for Adobe Flash Player" (http://helpx.adobe.com/security/products/flash-player/apsb14-26.html).</p>
29	
30	<p>S 27th: A U.S. SNS coordination service was targeted in a domain hijacking caused by a registry attack, leading to incidents in which certain messages and images were displayed on a large number of sites that used this service, including major domestic and international sites. See the following Gigya blog post for more details. "Regarding Today's Service Attack" (http://blog.gigya.com/regarding-todays-service-attack/).</p>

[Legend]



V Vulnerabilities



S Security Incidents



P Political and Social Situation



H History



O Other

*Dates are in Japan Standard Time

The website of a publisher in Japan was also compromised, causing visitors to be redirected to another website. In this incident, phishing emails were used to direct recipients to a login form misrepresented as being for the purpose of confirming member details, thus illegally acquiring information such as IDs and passwords. The compromise was caused due to the theft of an ID with cloud service administrator privileges.

■ Attacks Targeting the Internal Information of Corporations

During the current survey period, leaks of internal information through malware infections in corporate business systems continued to be discovered. In October, malware infections were uncovered at a number of retailers. It was announced that one of these was affected by the Backoff malware that targets POS systems, which has been used in multiple information leaks since last year^{*17}. US-CERT issued an alert regarding this malware in August^{*18}, but infections have continued to occur since then, indicating that it is still very active, and ongoing caution is required.

In November there was a large-scale attack on a major movie company in the U.S. In this incident, malware infections on internal PCs led to the leak of a large amount of internal data, such as unreleased movie data, passport information for movie-related persons, email data of employees, and other personal information. The fallout from this incident continued for some time after, including threats demanding that a certain movie stay unreleased, and the release of data leaked by an unknown party multiple times. Because the malware used included a function that disabled the specific security software used by the affected company, as well as a function that destroys the data on PCs, it is thought these attacks targeted the company in question and were executed after careful preparation. A number of security companies have identified that this malware was similar to the one used in mass simultaneous malware infections that occurred at a number of broadcasting stations and financial institutions in South Korea in 2013^{*19*20}.

In December, a power company in South Korea was attacked via emails with malware attached, leading to the leak of internal information. In this incident, the culprit demanded the shutdown of nuclear power stations as well as financial compensation, and released nuclear reactor plans and employee personal information on the Internet over a number of times.

■ Other

In the United States, the FBI announced they had shut down the Silk Road 2.0 crime site that dealt in illicit drugs and arrested its operator. In relation to this, the European Police Office (Europol) carried out a large-scale crackdown on crime sites operated on Tor in coordination with law enforcement agencies in 16 countries including the United States and European nations. In this operation, over 400 .onion domains that used Tor were seized. There were a large number of arrests, and Bitcoin currency, cash, a variety of drugs, and precious metals such as gold were confiscated.

In November, the Metropolitan and 19 other prefectural police departments established a joint investigation headquarters and carried out a nationwide simultaneous search due to the illegal operation of proxy servers using unlawfully obtained IDs and passwords. This resulted in the arrest of a number of people suspected of violating the Unauthorized Computer Access Law. In this incident, a number of people were also arrested on suspicion of copyright infringement (infringement of reproduction rights or business use of pirated copies) because pirated software was used on the servers operated^{*21}. Regarding the stolen IDs and passwords, it is thought that a known vulnerability in certain wireless LAN routers was exploited, and these details used to gain unauthorized access and conduct phishing attacks on major banks.

Also in November, the malware known as Regin received widespread attention. It is said a provider in Belgium was infected with this malware, which is designed to conceal itself within a network it has compromised and monitor communications. In

*17 International Dairy Queen, Inc. "Data Security Incident" (<http://www.dairyqueen.com/datasecurityincident/>).

*18 US-CERT, "Alert (TA14-212A) Backoff Point-of-Sale Malware" (<https://www.us-cert.gov/ncas/alerts/TA14-212A>).

*19 See "1.4.1 The 3.20 Cyber Attack in South Korea" in IIR Vol.19 (http://www.ij.ad.jp/en/company/development/iir/pdf/iir_vol19_EN.pdf) for more information about this incident.

*20 See the following Trend Micro security intelligence blog post for more information: "An Analysis of the 'Destructive' Malware Behind FBI Warnings" (<http://blog.trendmicro.com/trendlabs-security-intelligence/an-analysis-of-the-destructive-malware-behind-fbi-warnings/>).

*21 Association of Copyright for Computer Software "Nationwide Simultaneous Crackdown on Operators of Malicious Proxy Servers" (<http://www2.accsjp.or.jp/criminal/2014/post.php>) (in Japanese).

December Incidents

1	S 6th: An incident occurred in which the website of a publisher was compromised by an unknown party, and altered to redirect visitors to another website.
2	O 8th: The National Information Security Center held an intersectoral exercise for critical infrastructure.
3	"Summary of Intersectoral Exercise for Critical Infrastructure [2014 Annual Intersectoral Exercise]" (http://www.nisc.go.jp/active/infra/pdf/bunya_enshu2014gaiyou.pdf) (in Japanese).
4	V 9th: A vulnerability that could cause abnormal operations on a server or stop a service was discovered and fixed in a number of DNS software implementations such as BIND.
5	CERT/CC, "Vulnerability Note VU#264212 Recursive DNS resolver implementations may follow referrals infinitely" (https://www.kb.cert.org/vuls/id/264212).
6	V 10th: Microsoft published their Security Bulletin Summary for December 2014, and released a total of seven updates, including three critical updates such as MS14-080 and MS14-081, as well as four important updates.
7	"Microsoft Security Bulletin Summary for December 2014" (https://technet.microsoft.com/library/security/ms14-dec).
8	V 10th: A number of vulnerabilities in Adobe Flash Player that could allow arbitrary code execution were discovered and fixed.
9	"Security updates available for Adobe Flash Player" (http://helpx.adobe.com/security/products/flash-player/apsb14-27.html).
10	V 10th: A number of vulnerabilities in Adobe Reader and Acrobat that could allow arbitrary code execution were discovered and fixed.
10	"APSB14-28: Security updates available for Adobe Reader and Acrobat" (http://helpx.adobe.com/security/products/reader/apsb14-28.html).
11	S 15th: Korea Hydro & Nuclear Power was attacked using email with malware attached, leading to the leak of nuclear reactor plans and manuals, as well as the personal information of employees.
12	S 17th: ICANN announced that a number of systems had been accessed without authorization, and some user information had leaked.
13	See the following ICANN announcement for more details: "ICANN Targeted in Spear Phishing Attack Enhanced Security Measures Implemented" (https://www.icann.org/news/announcement-2-2014-12-16-en).
14	O 17th: The Federal Office for Information Security (BSI) of the German Federal Ministry of the Interior published its 2014 IT security white paper.
15	The white paper states that an incident occurred in which a steel mill was attacked, causing damage to the facilities. See "3.3.1 APT attack on industrial installations in Germany" in "The State of IT Security in Germany 2014" (https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/Securitysituation/IT-Security-Situation-in-Germany-2014_pdf.pdf?__blob=publicationFile) for more information about this incident.
16	O 18th: CODE BLUE, an international information security conference originating in Japan, was held over two days.
17	See the following official site for CODE BLUE for more information (http://www.codeblue.jp/2014/en/).
18	O 19th: The JPCERT Coordination Center issued an alert due to an increase in scans to TCP port 8080 from around December 5. This was scanning behavior targeting the NAS products provided by QNAP Systems, Inc.
19	"JPCERT/CC Alert 2014-12-19 Alert regarding increase in scans to TCP port 8080" (http://www.jpcert.or.jp/english/at/2014/at140055.html).
20	V 20th: A number of vulnerabilities in ntpd that could allow arbitrary code execution with ntpd privileges by sending specially-crafted packets were discovered and fixed.
21	CERT/CC, "Vulnerability Note VU#852879 NTP Project Network Time Protocol daemon (ntpd) contains multiple vulnerabilities (Updated)" (http://www.kb.cert.org/vuls/id/852879).
22	O 22nd: The Ministry of Internal Affairs and Communications made revisions to their "Guidelines for Unlocking Sim Lock," which is one of the initiatives aimed at lowering the cost of mobile services and encouraging more diversity. In these guidelines it is stated that all devices sold by carriers should support the unlocking of SIM cards in principle.
23	"Revisions to 'Guidelines for Unlocking SIM Lock'" (http://www.soumu.go.jp/menu_news/s-news/01kiban03_02000275.html) (in Japanese).
24	S 23rd: An incident occurred in which the website of the Internet Systems Consortium (ISC) was accessed without authorization, and altered to redirect visitors to a malware distribution site.
25	See the following blog post from U.S. security firm Cyphort that discovered the issue for more information. "Internet Systems Consortium's ISC.org infected" (http://www.cyphort.com/isc-org-infected/).
26	V 28th: At a security conference held in Germany, it was announced that the SS7 telecommunication standard used in public telephone networks has a vulnerability that could allow wiretapping.
27	See the following announcement for more information about this vulnerability. Laurent Ghigonis, Alexandre De Oliveira, "SS7map : mapping vulnerability of the international mobile roaming infrastructure" (http://media.ccc.de/browse/congress/2014/31c3_-_6531_-_en_-_saal_6_-_201412272300_-_ss7map_mapping_vulnerability_of_the_international_mobile_roaming_infrastructure_-_laurent_ghigonis_-_alexandre_de_oliveira.html#video).
28	V 29th: Google announced it had found a vulnerability with no fix available in Microsoft's Windows 8.1 Update that could lead to the elevation of privileges.
29	This information was disclosed at google-security-research (https://code.google.com/p/google-security-research/).
30	
31	

[Legend]

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

*Dates are in Japan Standard Time

addition to standard RAT functions such as password theft and network traffic monitoring, special functions such as a feature for monitoring mobile phone base stations have also been confirmed*22.

In December, the Federal Office for Information Security (BSI) of the German Federal Ministry of the Interior published an IT security white paper, which revealed that attacks on a steel mill had caused damage to production facilities due to their shutdown.

A number of large-scale DDoS attacks occurred during this period. In December, an unknown entity calling themselves the Lizard Squad attacked a number of online game services such as PSN and Xbox Live, as well as Tor. The attackers subsequently provided the attack platform used in this series of attacks as a DDoS attack tool, and this was actually used in DDoS attacks on a number of sites. In January a number of people thought to be members were arrested, but as attacks continue to be made by remaining members afterwards, ongoing vigilance is necessary.

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 October 1 and December 31, 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: attacks on bandwidth capacity*23, attacks on servers*24, and compound attacks (several types of attacks on a single target conducted at the same time).

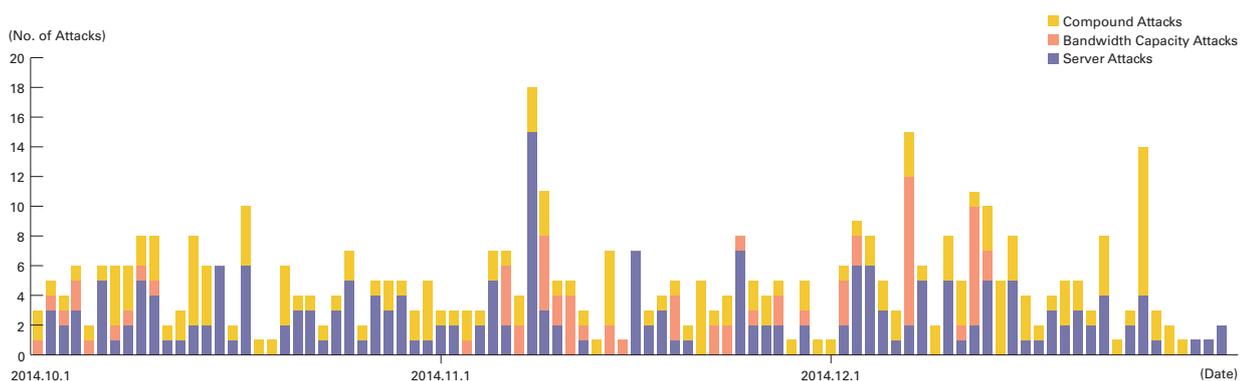


Figure 2: Trends in DDoS Attacks

*22 See the following Symantec Security Response Blog post for more information about this malware. "Regin: Top-tier espionage tool enables stealthy surveillance" (<http://www.symantec.com/connect/blogs/regin-top-tier-espionage-tool-enables-stealthy-surveillance>)

*23 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.

*24 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.

During the three months under study, IJ dealt with 381 DDoS attacks. This averages to 4.14 attacks per day, indicating an increase in the average daily number of attacks compared to our prior report. Server attacks accounted for 54.6% of all incidents, while compound attacks accounted for 26.9%, and bandwidth capacity attacks 18.6%. The largest attack observed during the period under study was classified as a compound attack, and resulted in 3.1 Gbps of bandwidth using up to 432,000 pps packets.

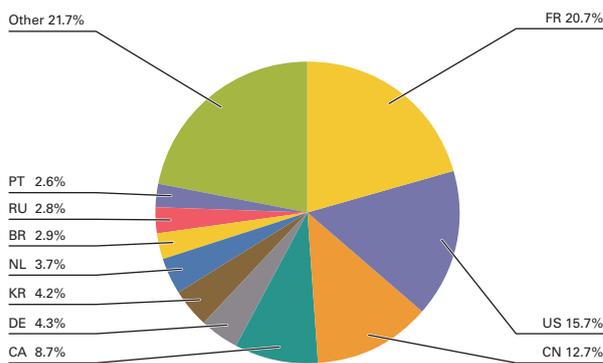
Of all attacks, 90.8% ended within 30 minutes of commencement, 9.2% lasted between 30 minutes and 24 hours, and none lasted over 24 hours. The longest sustained attack was a compound attack that lasted for six hours and 28 minutes.

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^{*25} and botnet^{*26} usage as the method for conducting DDoS attacks.

■ Backscatter Observations

Next we present our observations of DDoS attack backscatter using the honeypots^{*27} set up by the MITF, a malware activity observation project operated by IJ^{*28}. 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 October 1 and December 31, 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 80/TCP port used for Web services, accounting for 35.1% of the total during the target period. This was followed by 53/UDP used for DNS at 22.8%, so the top two ports accounted for 57.9% of the total. Attacks were also observed on 6667/TCP used for IRC (Internet Relay Chat) and 443/TCP used for HTTPS, as well as the typically unused 0/TCP, 34000/TCP, and 25565/TCP.

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

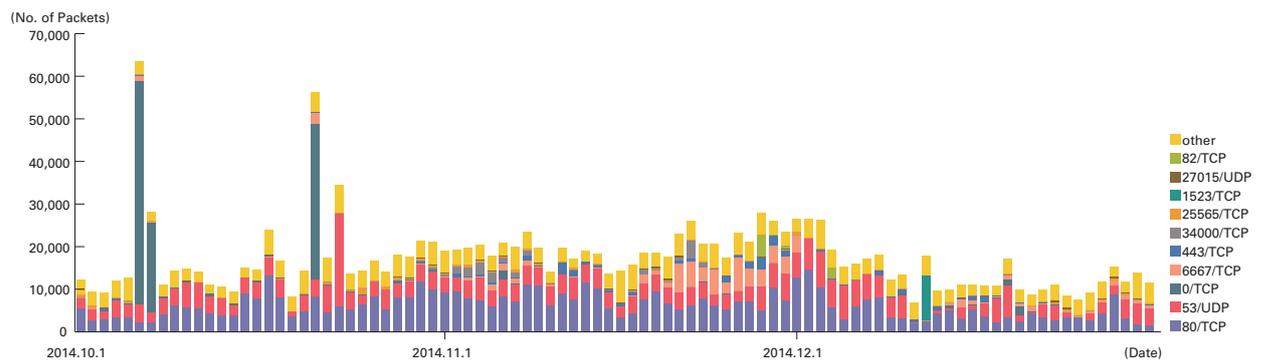


Figure 4: Observations of Backscatter Caused by DDoS Attacks (Observed Packets, Trends by Port)

*25 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.

*26 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.

*27 Honeypots established by the MITF, a malware activity observation project operated by IJ. See also “1.3.2 Malware Activities.”

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

The 53/UDP backscatter that has been prevalent in observations since February 2014 remains the second highest in number of packets observed by port in this survey. However, the daily average has increased from approximately 3,100 packets to approximately 3,900 packets.

Looking at the origin of backscatter thought to indicate IP addresses targeted by DDoS by country in Figure 3, France accounted for the largest ratio at 20.7%. The United States and China followed at 15.7% and 12.7%, respectively.

Looking at particularly large numbers of backscatter packets observed by targeted port, there were attacks on the Web servers (80/TCP) of a news site in Portugal between October 16 and November 23, and on a game-related site in Germany between November 17 and November 20. Attacks were also observed on a Spanish news site between October 14 and October 18. Between October 6 and October 7, as well as on October 21, many attacks were observed on 0/TCP. Backscatter from a wide range of IP addresses was received on many honeypots, and the purpose of the attacks is not known. Other observations included ongoing attacks from October 23 on 80/TCP, 6667/TCP, and 34000/TCP targeting IRC servers in France.

Notable DDoS attacks during the current survey period that were detected via IIJ's observations of backscatter included attacks on the site of a municipal authority in the U.S. city of Phoenix, Arizona on October 22. Attacks were also detected on Ukraine's election commission on October 26, and on the site for a political party in France on November 29.

1.3.2 Malware Activities

Here, we will discuss the results of the observations of the MITF*²⁹, a malware activity observation project operated by IIJ. The MITF uses honeypots*³⁰ 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.

■ Status of Random Communications

Figure 5 shows the distribution of sender's IP addresses by country for communications coming into the honeypots between October 1 and December 31, 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

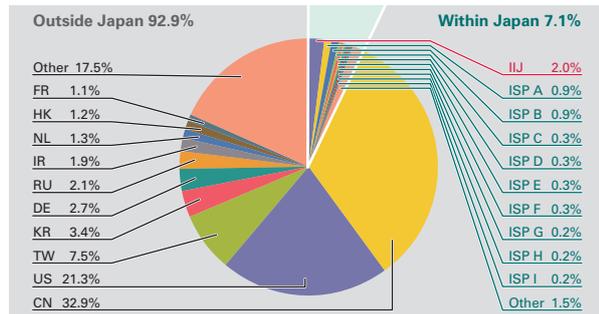


Figure 5: Sender Distribution (by Country, Entire Period under Study)

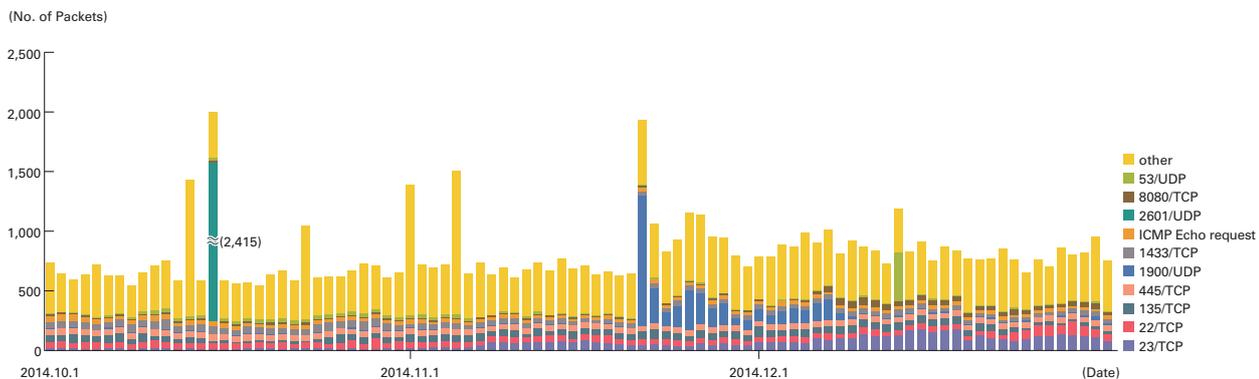


Figure 6: Communications Arriving at Honeypots (by Date, by Target Port, per Honeypot)

*29 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.

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

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 period, there was an increase in 1900/UDP communications used in UPnP's SSDP between late November and early December. These packets involved repeated m-search requests within a short time period. It is thought these were attempted SSDP reflection attacks (a type of DDoS attack) with the sender's IP addresses spoofed, which means these source addresses in fact indicated the targets of the attacks*³¹. The National Police Agency issued an alert regarding these incidents in October*³².

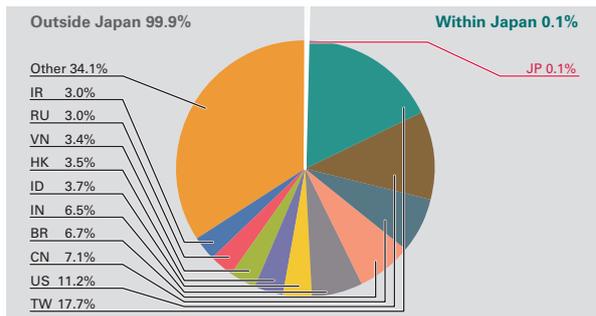


Figure 7: Distribution of the Number of Malware Specimens Acquired

From November there was an increase in communications targeting telnet (23/TCP). Upon investigation, we learned that this was mainly received from IP addresses allocated to China and Japan. From December 5, there was an increase in communications targeting 8080/TCP. We believe these were attacks on the Shellshock vulnerability in QNAP

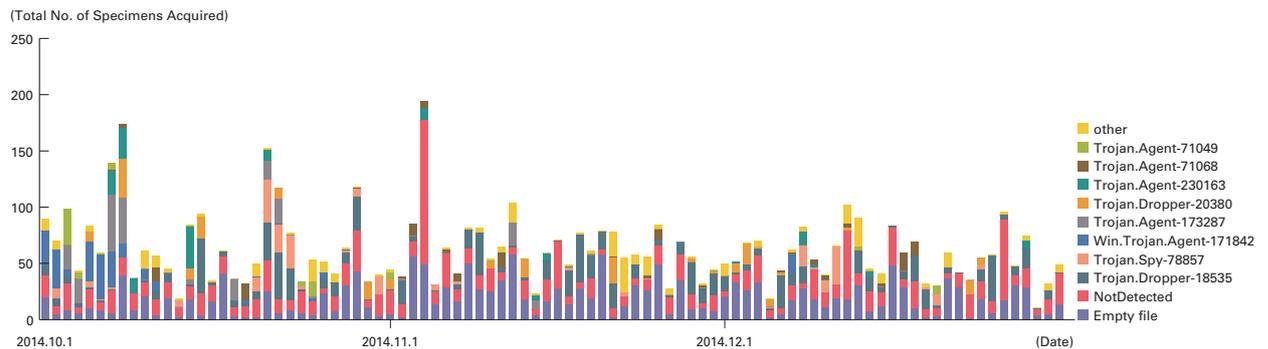


Figure 8: Trends in the Total Number of Malware Specimens Acquired (Excluding Conficker)

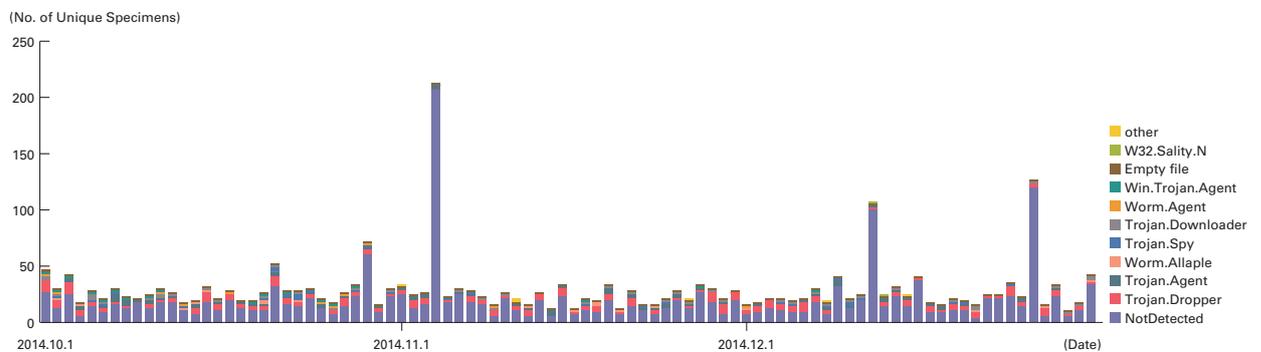


Figure 9: Trends in the Number of Unique Specimens (Excluding Conficker)

*31 MITF honeypots have SSDP disabled, so do not assist the attacks.

*32 "Alert regarding SSDP reflector attacks using UPnP compatible network devices as stepping stones" (<http://www.npa.go.jp/cyberpolice/detect/pdf/20141017.pdf>) (in Japanese).

brand NAS products^{*33}, etc. On December 13, there was a large volume of communications targeting 53/UDP. Most of these packets were due to incoming DNS amp attack attempts^{*34}. Because the source was spoofed to an IP address assigned to a telecommunications company in Germany, it is possible that this IP address was the target. The queries contained A record resolution requests for an existing domain, but approximately 250 A records were set to that FQDN. This resulted in a DNS amp attack, as using a spoofed source to perform name resolution for a domain set with a large number of A records causes responses to be amplified. On October 15, communications targeting 2601/UDP were made to the IP address of a specific honeypot from an IP address allocated to China. Upon investigating these communications, we found that random data with a length between 50 bytes and around 800 bytes had been sent repeatedly within a short space of time, but the purpose is not known.

■ Malware Network Activity

Figure 7 shows the distribution of the specimen^{*35} 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, while the number of unique specimens is the number of specimen variants categorized according to their digest of a hash function^{*36}. 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.

On average, 63 specimens were acquired per day during the period under study, representing 17 different malware. After investigating the undetected specimens more closely, they included worms observed from IP addresses allocated to countries such as India, the United States, and Taiwan. 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 93.4% of malware specimens acquired were worms, and 6.6% were downloaders. In addition, the MITF confirmed the presence of 106 botnet C&C servers^{*37} and 7 malware distribution sites. The number of botnet C&C servers rose dramatically, but this was due to the appearance of a specimen that used a DGA (Domain Generation Algorithm) during the current survey period.

■ Conficker Activity

Including Conficker, an average of 16,343 specimens were acquired per day during the period covered by this report, representing 557 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 dropped considerably, at approximately 36% lower than the previous survey period. Unique specimens were also down by about 17%. According to the observations of the Conficker Working Group^{*38}, as of January 13, 2015, a total of 890,845 unique IP addresses are infected^{*39}. This indicates a drop to about 28% of the 3.2 million PCs observed in November 2011, but it demonstrates that infections are still widespread.

*33 It was also reported in JPCERT/CC observations that scanning and attacks increased on the same day. "Alert regarding increase in scans to TCP port 8080" (<http://www.jpcert.or.jp/english/at/2014/at140055.html>).

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

*35 This indicates the malware acquired by honeypots.

*36 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.

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

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

*39 For some reason Conficker Working Group data appears to be missing for about a month from mid-December 2014, so we have cited data for January 13, 2015 that should not be affected.

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 October 1 and December 31, 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.

Japan was the source for 30.8% of attacks observed, while the United States and China accounted for 13.5% and 13.4%, respectively, with other countries following in order. There was a decrease in the number of SQL injection attacks on Web servers compared to the previous report, but this is due to a significant spike in attacks originating from China in the last survey period, so there was no change in overall detection trends.

During this period, large-scale attacks from a specific attack source in France directed at specific targets took place on October 12. Attacks on these targets were confirmed from specific attack sources in Russia on October 10 and October 16, and from specific attack sources in Indonesia on December 18. On October 6, attacks were made from a number of sources in Europe and the United States directed at specific targets. On October 7, attacks originating in the Czech Republic were made on specific targets. Attacks from a specific attack source in Japan directed at specific targets also took place on October 23. On November 24, there were attacks from a specific attack source in China directed at specific targets. 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, as these attacks are ongoing, they continue to require attention.

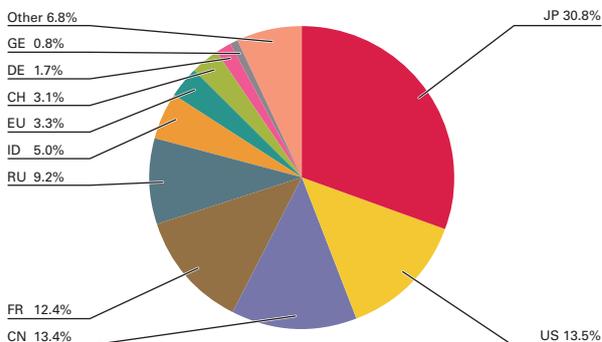


Figure 10: Distribution of SQL Injection Attacks by Source

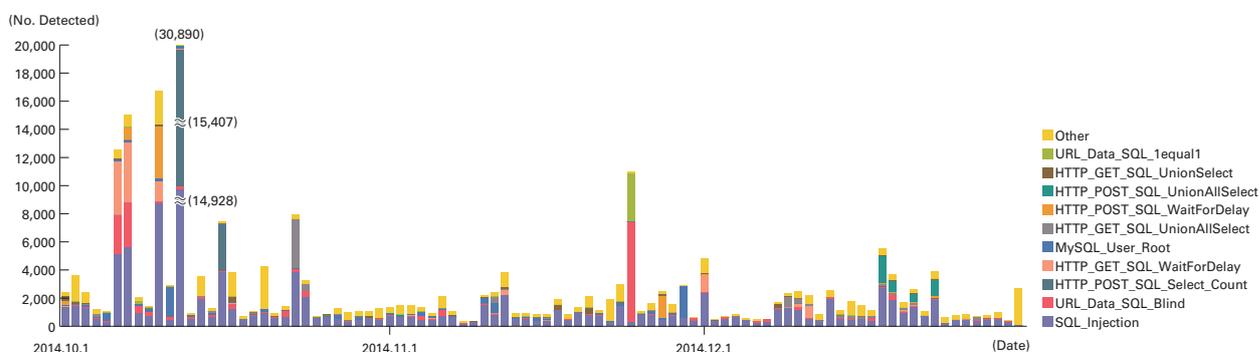


Figure 11: Trends in SQL Injection Attacks (by Day, by Attack Type)

*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.

The number of drive-by download attacks observed between October and December 2014 fell to about a third of those in the period from July to September (Figure 12). In particular, there were many days in December in which no attacks were observed. Anglar continued to account for many attacks as in the previous survey, making up the majority along with Fiesta, which saw a sudden rise in this survey period. Like most other exploit kits, Fiesta is equipped with functions that exploit vulnerabilities in Internet Explorer and its plug-ins, as well as Flash, Java, and Silverlight. A recent trend in many exploit kits such as Anglar and Fiesta is that functions for exploiting a number of comparatively new Flash vulnerabilities (including CVE-2014-8439, CVE-2014-0515, and CVE-2014-0497) have been added at a rapid pace*⁴². Conversely, we no longer see as many functions exploiting Java vulnerabilities that were previously often targeted. We believe this is because of improvements made to Java-related security functions in popular browser environments. In 2014, a function that blocks the automatic execution of older versions of Java was added to Internet Explorer, and the automatic execution of Java was also completely blocked in Chrome and Firefox.

Most websites altered to redirect visitors were small-scale sites that are less popular, and we observed intermittently that individual sites were used as redirect sites for periods between a few days and approximately two months. With regard to trends in content, there were many sites that introduce adult-oriented video content, and websites for idol groups and design businesses among others were also altered.

Overall, it is estimated that the incidence rate for drive-by downloads is declining sharply. However, there is always the possibility of a sudden change in trends such as this depending on the intentions of attackers. It is recommended that all parties 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 (Flash Player in particular).

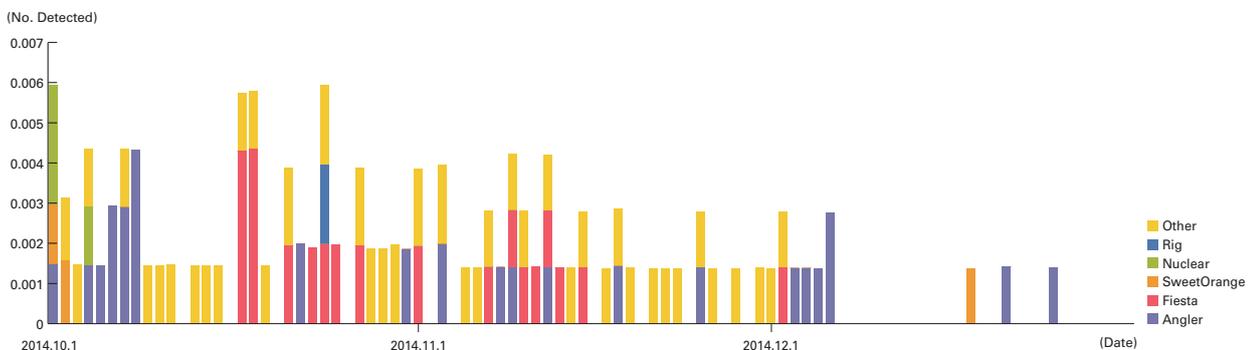


Figure 12: Drive-By Download Frequency When Viewing Websites

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

*⁴² Details of the vulnerabilities targeted by each exploit kit are summarized in the "ExploitPackTable_2014" (<https://docs.google.com/spreadsheets/ccc?key=0AjvsQV3iSLa1dE9EVGhjeUhvQTNReko3c2xhTmphLUE>)

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, including an overview of countermeasures against the alteration of domain name registration information, and discussion of the `openioc_scan` plug-in used to scan for threats that lurk in the memory of a computer. We also take a look at ID management technology.

1.4.1 Countermeasures Against the Alteration of Domain Name Registration Information

■ The Threat of Altered Domain Name Registration Information

Domain names play an important role in the Internet world. For example, when a client accesses `www.example.com`, it sends a query to the authoritative DNS server for the `example.com` zone that the server is associated with to obtain that server's IP address (Figure 13 left). This series of exchanges is called name resolution.

However, there are attacks that cause a server prepared by an attacker to be accessed instead of the legitimate server as a result of name resolution. These attacks are called domain hijackings. The attack techniques used include compromising the DNS server of the domain owner and rewriting records, making a recursive DNS server cache a malicious IP address using DNS cache poisoning, and routing specific traffic to a network other than the intended one by advertising incorrect information to BGP.

The impact of attack techniques such as these is diminishing due to countermeasures implemented by ISPs and service providers. However, a number of domain hijackings using a different attack technique occurred in September 2014 and beyond, affecting the domains of companies in Japan. This attack technique involves altering the registration information for a domain name registered to a registry so that it responds with the IP address of a server prepared by the attacker when name resolution is performed. As a result, even if the correct domain name is entered, the server prepared by the attacker is accessed, potentially exposing users to harm such as phishing attempts or malware infections. One characteristic of this technique is that servers are affected by attacks even if their security level is high, because fake information is registered on the authoritative DNS server for the top-level domain (TLD)^{*43} (Figure 13 right). Furthermore, because the fake information

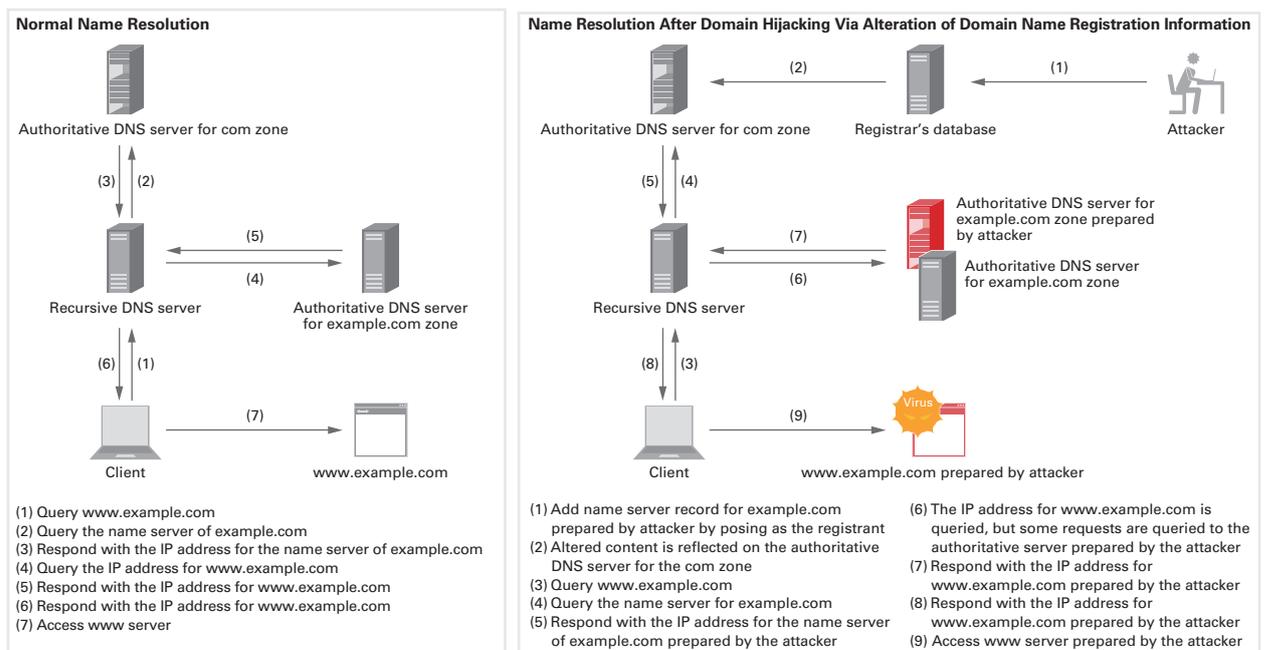


Figure 13: Name Resolution of `www.example.com` (DNS protocols partially omitted)

*43 Indicates the right-most label of domain names delimited by a dot (`.`), such as `jp` and `com` in `www.example.jp` and `mail.example.com`, respectively.

only has to be registered to an authoritative DNS server upstream of the domain the server belongs to, when the domain hierarchy is deep, the fake information may be registered to an authoritative DNS server of a second-level domain (SLD) rather than a TLD.

In the incidents that occurred between September and October 2014, attacks that caused Web page visitors to download malware were made on “.com” domains owned by a number of companies in Japan^{*44*45}. Additionally, in November 2014 there were incidents affecting major domestic and international newspaper publishers and news sites using an SNS coordination service provided by an overseas company. In these incidents, Web page visitors were served with malicious JavaScript files that caused the crest of the Syrian Electronic Army (SEA) to be displayed.

In the incidents in November 2014, all NS records were altered to those of the attacker, forcing the malicious JavaScript to be always served. However, in the incidents affecting the domains of domestic companies between September and October, the attacker’s NS records were added without deleting the legitimate NS records. For this reason, only some users were affected.

Many attacks using techniques similar to these have already occurred overseas. For example, in 2013 a large number of registries were attacked, and incidents occurred on almost a monthly basis^{*46}.

■ Techniques for Altering Domain Name Registration Information

Information regarding domains is registered in the databases of organizations called registries, which operate and manage TLDs. To register a domain, a domain registration application is sent to a registry using the services of an intermediary called a registrar (or a reseller of the registrar’s services). Domain registration information is passed on in the order shown by the arrows in Figure 14, and registered to the registry’s authoritative DNS server and WHOIS server.

Consequently, attackers alter domain name registration information by attacking any of the points below.

- (1) Registrar data is altered by stealing the identity of the registrant or reseller.
- (2) Registry data is altered by stealing the identity of the registrar.
- (3) Data is altered using a vulnerability on the system of the registry or registrar.

Identity theft is carried out based on information gained through methods such as phishing, malware infections, or social engineering targeting the domain registrant. Leaked account information for registries or registrars and list-based attacks are also potential causes.

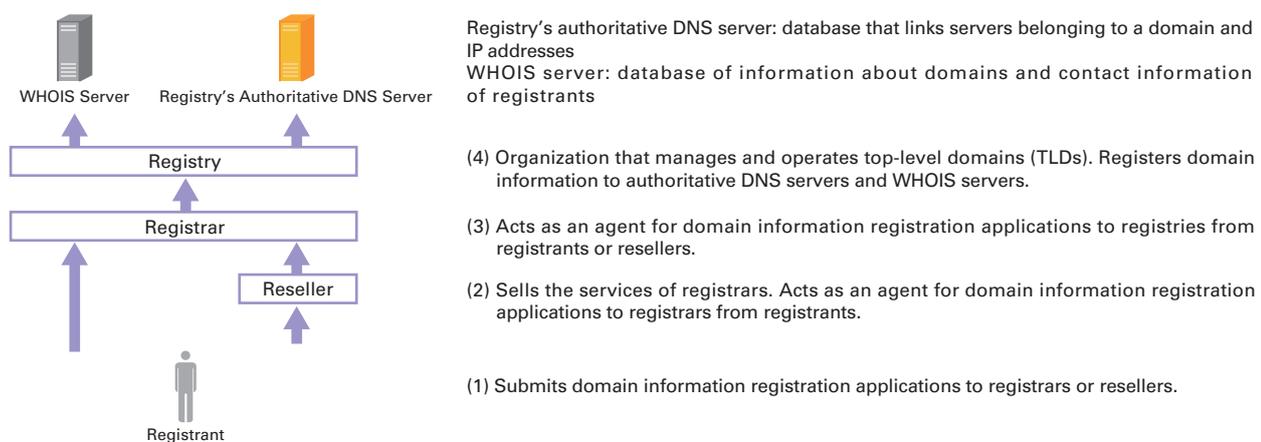


Figure 14: Domain Information Registration Process

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

*45 IJ has not observed cases in which malware infections took place.

*46 Japan Registry Services (JPRS), “Additional Reference: Regarding domain name hijackings caused by unauthorized rewriting of registered information and their countermeasures” (<http://jprs.jp/tech/security/2014-11-05-unauthorized-update-of-registration-information.pdf>) (in Japanese).

The data altered is the name servers registered to the registry's WHOIS server and the NS records and glue records registered to the authoritative DNS server.

■ Countermeasures

Here we explain countermeasures against the alteration of domain name registration information.

■ Registries/Registrars/Resellers

The most fundamental security measure is to ensure that fixes and patches for vulnerabilities in OSes, Web applications, and APIs are applied. Measures such as applying patches and installing antivirus software on clients within the organization must also be taken.

Next you should contemplate improving functions for authenticating user accounts to protect against identity theft. It is likely that many systems perform authentication using passwords, so it is necessary to set complicated passwords and prevent reuse of the same password. If possible, also consider implementing two-step authentication or authentication based on client certificates.

Additionally, look into implementing functions for monitoring and sending notification of user activity. It is possible to detect actions not intended by the user by sending notification of account logins and changes to registration information. Furthermore, if there is a history of this, users can confirm when unauthorized access began themselves.

To protect against actual attacks, it is necessary to implement systems for detecting and dealing with attacks. Attacks are detected based on security devices, system load, or logs, and blocked using a firewall, etc., as appropriate.

Finally, consider providing a Registry Lock service. This is a function that restricts changes to a registration information at a registry, making it possible to prevent unintended changes to information because additional authentication is required of the registrant when changing registration information.

There have also been cases in which services for resetting email addresses or passwords by fax have been exploited at overseas registrars. If you provide a similar service, it would be advisable to review the need for this service and the process for identity confirmation.

■ Registrants

It is important to select a registry/registrar that implements measures such as those detailed above. Proactively use features provided by registrars such as improved authentication, notification functions, and Registry Lock. Also apply OS patches and use the latest versions of software. Additionally, install antivirus software. You should also never reuse the same passwords for registrar accounts.

When using a notification function, configure settings so that notification emails are not detected as spam. When doing this, the contact email address should be on a domain other than one you have registered. This enables you to contact the registrar, even if domain name registration information is altered. If there is an option for not displaying registrant information in the public WHOIS information, consider enabling this.

Furthermore, you may be able to detect the alteration of registration information earlier by periodically confirming that information registered to the WHOIS server and registry's authoritative DNS server has not been altered. Check the name server on the WHOIS server, and the domain's NS records and glue records on the registry's authoritative DNS server.

However, because there is no standardized tool for monitoring this information, it is necessary to create your own. We have confirmed a number of monitoring tools released on voluntary basis, but it is necessary to confirm their behavior sufficiently before using them, since some only monitor specific records or do not send queries directly to a registrar's authoritative DNS server.

Furthermore, when performing monitoring, do not send more queries than necessary to a registry's authoritative DNS server or WHOIS server. If you send excessive queries, they may be detected as an attack on the registry, causing them to be restricted or blocked.

■ General Users

It is almost impossible for general users to determine whether or not domain name registration information has been altered. Consequently, apply OS patches, use the latest versions of software, and install antivirus software to prevent vulnerabilities being exploited when you access a server prepared by an attacker.

Additionally, because it is not possible to steal the SSL/TLS private keys of a server by altering domain name registration information, attackers cannot provide services that carry out SSL/TLS encryption. For this reason, when you access a service normally provided via HTTPS using HTTP, or encounter SSL/TLS errors, it is possible that domain name registration information has been altered, so it would be advisable to stop using that service.

■ Summary

At this point in time the most effective countermeasure is Registry Lock, but some registries do not provide this function. Furthermore, by implementing the countermeasures detailed here regardless of whether Registry Lock is used, you can prevent the alteration of domain name registration information, or detect it at an early stage, just in case Registry Lock is compromised.

The larger the number of users that a domain provides services to, the greater the impact, so take this opportunity to implement countermeasures now.

1.4.2 The openioc_scan Plug-in That Scans for Threats Lurking in Device Memory

An IOC (indicator of compromise) is an artifact left on a network or computer that indicates a threat of malware infection or intrusion. By defining IOCs based on the results of malware or forensic analysis, it is possible to detect the same threat swiftly the next time it occurs. In this section, we introduce a tool IIJ has implemented that scans for IOCs in computer memory, and discuss a study regarding using it with generic IOCs.

■ Background to Implementation

The term IOC does not refer to a specific format or implementation. At a site called IOC Bucket^{*47}, IOCs are shared in a number of formats. Figure 15 shows the number and ratio of IOC formats shared on this site (as of January 8, 2015).

From this figure, we can see that OpenIOC^{*48} accounts for over three quarters of all formats, making it the most prevalent major format. As mentioned in a previous IIJ SECT blog article^{*49}, there are free tools such as IOC Finder and Redline that perform scans based on OpenIOC definitions. IOC Finder is a tool for scanning systems as they are running, enabling an effective live response. Meanwhile, Redline scans saved memory images offline. In the latter case, being offline makes it possible to disable the concealment of information by malware, and define the characteristics of strings in memory as well as code after unpacking, so you can make definitions based on malware functions.

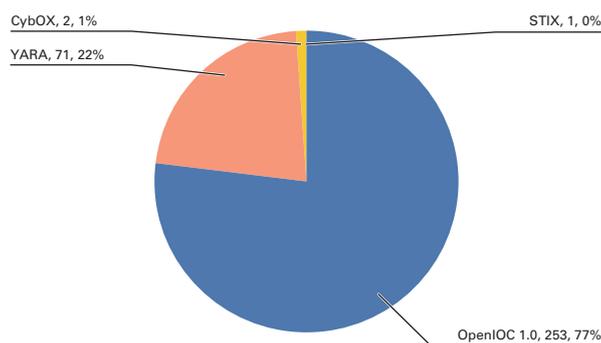


Figure 15: Number and Ratio of IOC Formats Shared on IOC Bucket

Up until now, IIJ has conducted evaluations for the early detection of malware using OpenIOC and Redline^{*50}. However, Redline is a closed source tool, so there was a problem in that we could not extend its functions or fix bugs ourselves. Consequently, we implemented a plug-in for the open source Volatility Framework^{*51} tool, called openioc_scan^{*52}.

*47 An IOC sharing site run by volunteers. As of January 8, 2015, 327 IOCs have been shared (<https://www.iocbucket.com/search>).

*48 A standard promoted by Mandiant. IOCs are written in XML format (<http://openioc.org/>).

*49 This article describes examples utilizing the OpenIOC free tool. "Defining and detecting traces of the presence of threats using OpenIOC" (<https://sect.ij.ad.jp/d/2012/02/278431.html>) (in Japanese).

*50 At the SANS DFIR Summit 2013 the year before last, we gave a presentation regarding techniques for defining and detecting volatile IOCs (https://digital-forensics.sans.org/summit-archives/DFIR_Summit/Volatile-IOCs-for-Fast-Incident-Response-Haruyama.pdf).

*51 An open source memory forensic tool for which a variety of plug-ins are made available by volunteers (<https://github.com/volatilityfoundation/volatility>).

*52 The latest version is available from the following link (<http://takahiroharuyama.github.io/blog/2014/08/15/fast-malware-triage-using-openioc-scan-volatility-plugin/>).

■ openioc_scan

Here we will describe how `openioc_scan` is used. Analysis using `openioc_scan` can only be performed on Windows OSes from Vista onward^{*53}, as Linux and Mac OS X are not currently supported. Additionally, the three Python packages `lxml`^{*54}, `ioc_writer`^{*55}, and `colorama`^{*56} are required to execute it. IOCs must also be defined in advance. To support regular expressions and the parameters described below, `openioc_scan` uses the OpenIOC 1.1 format. Currently, the only free tool capable of OpenIOC 1.1 definitions is `PyIOCe`^{*57}, which we will use to define IOCs in the following description.

Figure 16 shows part of the `PyIOCe` screen. Users define the terms (items) for volatility on this screen. In the figure, the two terms `ProcessItem/ParentProcessName` and `ProcessItem/name` are combined using AND/OR logic and defined. In addition to being able to specify NOT (negation) as shown in the figure, it is also possible to differentiate between upper-case and lower-case characters, and specify matching^{*58}.

Once IOCs are defined, the folder where these definition files are located is specified using the `ioc_dir` option, and `openioc_scan` is executed as shown in Figure 17. In principle, if combining the evaluation results for each term using AND/OR results in a true match, the term that is a true match for that IOC is displayed in a different color. In the figure, we can see that `svchost.exe` with the process ID 2204 is the process that matched the criteria. Incidentally, unlike Redline, Volatility Framework does not cache the analysis results for memory images. However, `openioc_scan` caches the information required for the evaluation of each term on a case-by-case basis, so it is possible to perform processing at high speed when evaluating the same term a second and subsequent time^{*59}.



Figure 16: Example of IOC Definition Using PyIOCe

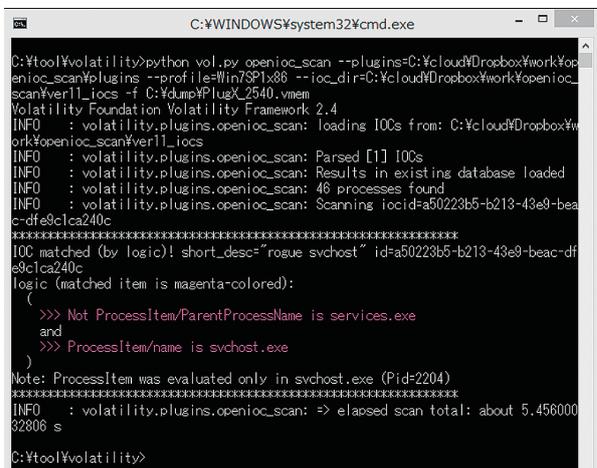


Figure 17: `openioc_scan` Execution

Table 1 shows a list of terms that `openioc_scan` supports. Regarding `ProcessItem` and `DriverItem`, the processes and drivers to be scanned are evaluated individually to check if they are a true or false match with IOC definitions. For example, the IOCs in the previous example are evaluated within a single process. It is also possible to make definitions that combine terms from different categories. However, this may dramatically decrease performance depending on the combination^{*60}, so for definitions it would be better to define terms simply in a single category whenever possible.

Previously, it was stated that in principle, if combining the evaluation results for each term using AND/OR results in a true match, the term that is a true match for that IOC is displayed in a different color. In fact, there is a method for displaying IOCs even when the final result of AND/OR for a term is not true.

OpenIOC 1.1 uses the concept of parameters to enable metadata to be defined for each term. In `openioc_scan`, scoring-based evaluation is supported using these

*53 It is actually also possible to scan memory images of XP and 2003 for items other than communications information, but this is not supported as sufficient testing has not been carried out.

*54 A package for parsing XML (<https://pypi.python.org/pypi/lxml/3.2.1>).

*55 A package for parsing OpenIOC 1.1 items (https://github.com/mandiant/ioc_writer)

*56 A package for changing the color of the text output in the console (<https://pypi.python.org/pypi/colorama>).

*57 An open source tool by Sean Gillespie (<https://github.com/yahoo/PyIOCe>).

*58 It is possible to specify `is/contains/matches/starts-with/ends-with/greater-than/less-than` based on factors such as whether the object to be scanned is a character string or numerical value. Regular expressions are specified for matches.

*59 Only information related to terms that take time are cached, such as the extraction of character strings. If the type of term is the same, it is processed at high speed even if the values are changed.

*60 For example, `ProcessItem` and `DriverItem` often involve iterative processing, so when making definitions that combine terms belonging to each of these categories, execution tends to take longer.

parameters. For example, when evaluating the IOC in the previous example, even if only one of the terms is true, when a score defined as a parameter exceeds the threshold, the fact there was a match to the IOC based on the score is shown. Specifically, when a score set to the parameters exceeds an integer value total of 100, the corresponding IOC is displayed as shown in Figure 18. In `openioc_scan`, parameters such as `detail` and `note` are also supported^{*61}.

■ Study Regarding Generic IOCs

Generally, IOCs are used to detect known threats. The main reason for this is that the majority of previous IOCs are defined with information such as the MD5 hash signature of malware or the URL of a C&C server, which would be hard to reuse to detect unknown threats. Consequently, IIJ carried out a study of generic IOCs for `openioc_scan` that are independent of specific malware or incidents.

Table 2 summarizes the evaluation results. With generic definitions you cannot completely avoid false positives, so it is not possible for just anyone to detect unknown threats using the IOCs studied here. However, from the perspective of assigning

Table 1: Terms Supported by `openioc_scan`

Term Category	Term Type
ProcessItem	Process name, path name, argument, parent process name, DLL path name, presence of DKOM ^{*62} , presence of code injection, used API name, character string, handle name, network connection information, hooked API name, effective privilege type
RegistryItem	History of executable files accessed by OS (ShimCache ^{*63})
ServiceItem	Service name, descriptive name, command line
DriverItem	Driver name, used API name, character string, IRP function table ^{*64} hooks, callback function type, timer function presence
HookItem	Hooked SSDT entries ^{*65}
FileItem	Metadata type of file based on MFT entries ^{*66}

```

*****
IOC matched (by score)! short_desc="rogue svchost" id=a50223b5-b21
logic (matched item is magenta-colored):
(
  Not ProcessItem/ParentProcessName is services.exe
  and
  >>> ProcessItem/name is svchost.exe (score=100;)
)
Note: ProcessItem was evaluated only in svchost.exe (Pid=752)
*****

```

Figure 18: Example of Parameter Usage

Table 2: Results of Study Regarding Generic IOCs

IOC Definition	Study Details	Limitations
Abnormal execution path	Detects executable files that contain folders not normally used often in their path, such as <code>Recycle.Bin</code> or <code>Users\Public</code>	Effective against processes that are running, but generates many false positives when used on access history
Web injection	Detects processes with all <code>HttpSendRequest</code> API hooked	Fails to detect some malware depending on the type of hooking carried out ^{*67}
Code injection	Detects code-injected processes based on their memory space characteristics or the API used	Cannot detect wow64 processes ^{*68}
Position Independent Code (PIC) ^{*69}	Detects things such as code sequences of access to <code>PEB</code> ^{*70} and <code>GetPC</code> ^{*71} , and immediate values of API name hashes	False positives occur for <code>GetPC</code> , and searching for API name hash values creates significant load
UAC prompt bypass	Detects code sequences that inhibit UAC prompts	Only one technique is defined, and there are other techniques for inhibiting this
Storage of data in a special area of NTFS	Detects processes and drivers that read and write to the NTFS Extended Attributes ^{*72}	False positives occur when evaluating drivers
Lateral movement of targeted attacks	Detects attack tools used for lateral movement	Relies on metadata such as file names, so generic definitions are difficult

*61 The `detail` parameter displays a character string in its entirety when a partial match is found. Specify "on" as the value. The `note` parameter is used when you want to add a comment for each term. Detailed configuration methods are explained on the following page (<http://takahiroharuyama.github.io/blog/2014/10/24/openioc-parameters-used-by-openioc-scan/>).

*62 An abbreviation of Direct Kernel Object Manipulation. In this case, it indicates processes hidden by altering the linked list of processes retained by the kernel.

*63 A registry value found in `HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\AppCompatibility\AppCompatCache` or `HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\AppCompatCache\AppCompatCache`.

*64 IRP is an abbreviation of I/O Request Packets. The function table for controlling the reading and writing of the buffer set for each driver is called the IRP function table.

*65 An abbreviation of System Service Descriptor Table. SSDT entries are system call names.

*66 MFT is an abbreviation of Master File Table. MFT entries are the metadata created for each file and folder managed using an NTFS file system.

*67 With regard to inline hooks, Volatility Framework only checks the first three instructions, so when a hook is embedded in the fourth instruction or later it cannot be detected.

*68 Due to the limitations of the virtual address translation system of Volatility Framework.

*69 Code that operates regardless of the execution context, such as shellcode.

*70 An abbreviation of Process Environment Block. Information required to access the API used.

*71 An abbreviation of Get Program Counter. Code for acquiring the current execution point.

*72 A data space used for compatibility with OS/2 applications.

priority to conduct analysis in further detail (triage), we feel that they can be utilized to a certain extent. Meanwhile, we also saw cases in which detection failed due to the limitations of the functions in Volatility Framework itself.

■ Summary

In this section we introduced `openioc_scan`, and presented the results of our study regarding generic IOCs using this. The study identified limitations derived from the functions in Volatility Framework, but as this tool is open source unlike Redline, each user is able to freely fix and extend these functions.

It seems there are still few analysts that utilize IOCs at all, let alone apply the concept of scanning memory images for IOCs. It is easy to predict that the number of computers subject to analysis in incident responses, as well as memory and disk drive capacity, will continue to increase steadily in the future. Experienced analysts may already feel that existing techniques have reached their limits. Try using `openioc_scan` to detect threats and perform triage, enabling more efficient incident response.

1.4.3 ID Management Technology

In 2014, the never ending stream of unauthorized login incidents thought to use lists of user IDs and passwords leaked in other incidents and accidents continued. As a result, there is growing awareness of the dangers of authentication methods that use only IDs and passwords, and new ID management technology that also incorporates other authentication methods has attracted attention. In this section, we examine this ID (identity) management technology. ID management technology involves more than technology related to the lifecycle, such as the issuing, usage, and disposal of IDs. It also encompasses an extremely diverse range of approaches, such as user authentication, issuing and usage of credentials, management and circulation of attribute information, access control for various resources, and delegation of authority, as well as technology for coordinating this information between different entities. Consequently, there are a wide variety of specifications related to the management and coordination of IDs used on the Internet. Because each is developed based on different concepts, it is necessary to select the one most appropriate for the environment it will be used in. However, due to the different technical terms used and the increasingly broad range that specifications cover, this is unfortunately becoming an extremely difficult area to understand. In this section, we aim to assist the implementation of technology by avoiding discussion of specific specifications whenever possible, and examining general approaches to ID management technology.

■ Approaches to IDs - Entities and Identifiers

When it comes to identity, there are an extremely large number of approaches, definitions, and concepts, and this is one of the factors that causes confusion among technical experts and users. For that reason, in this report we will take the simple approach of treating IDs as identifiers.

Entities in the real world are tied to entities in the digital world. One example of an entity in the real world is a user that is trying to use some kind of service, but as represented by the IoT (Internet of Things), devices that actively seek to connect to other nodes also sometimes have IDs. When this happens, a unique identifier is assigned to identify entities in the digital world. IDs can be thought of as the members of the identifier space determined for each individual realm. ID uniqueness means that different IDs are assigned to different entities in the corresponding realm.

It is typical for entities in the real world to have different IDs in a number of realms. Furthermore, entities in the real world are often associated with multiple IDs in the same realm. However, in the digital world, different IDs are recognized as different entities. For example, the real-world entity Mr. A may have the email addresses `a@aaa.example` and `a@bbb.example`, possessing IDs in the two different realms, `aaa.example` and `bbb.example`. Mr. A may also have different email addresses in the same domain in some cases, such as `a1@aaa.example` in addition to `a@aaa.example` in the `aaa.example` realm. This demonstrates that the same entity in the real world can have multiple IDs assigned in different realms in the digital world.

■ Tokens/Credentials and Authentication

As mentioned previously, using IDs in the digital world enables us to identify different entities. However, IDs are public information, and it would pose a problem if simply anyone could claim they were a particular entity. Consequently, an authentication system is required for verifying an entity that is requesting access really is the entity that ID was assigned to. For this, private information such as a password that matches the corresponding ID is needed.

Here, in line with the NIST SP800-63^{*73} definitions, we will explain credentials^{*74} and tokens separately. Tokens indicate information that should be kept secret held by users assigned the corresponding ID. They fall into the categories of “something you know,” “something you have,” or “something you are.” Examples of tokens include passwords or the private keys used in public key cryptosystems, physical media such as IC cards or dongles, and body characteristics used for biometric authentication such as fingerprints or irises. Meanwhile, credentials are public information that indicate associations between tokens and IDs. In some cases credentials are signed by a trusted entity, and serve as information that a third party can verify the validity of. For password-type tokens, IDs can be thought of as credentials. Considering public key cryptosystem authentication methods used with protocols such as SSL/TLS, the X.509 certificate^{*75} is the credential, and the private key that matches the public key included in the certificate corresponds to the token. SSL/TLS server certificates also include the FQDN, and this can be thought of as an ID. Bitcoin addresses can be regarded as IDs, but because the address itself is a public key, and transaction signature verification can be performed using only the address, you can also regard them as credentials.

An authentication method using multiple [token/credential] pairs instead of a single pair to authenticate an entity with a certain ID assigned is called multi-factor authentication. The most widely-known method is parallel multi-factor authentication, in which independent tokens are used for each authentication method. Meanwhile, some opt for a cascading approach to multi-factor authentication. For example, when using a private key as a token in public key cryptosystems, the private key file is generally encrypted, so the password must be entered to decrypt it. The use of both the password and private key as tokens at this time can be regarded as cascading multi-factor authentication. The same applies to IC cards and PIN numbers. Additionally, for hardware tokens that fall into the “something you have” token category, a one-time password that is shown on the physical media and periodically updated is typically presented as a token, and used as a part of multi-factor authentication. For network environments, a one-time password system using Lamport’s hash chains^{*76} is known, and has been drawn up in specifications^{*77*78} such as S/Key. One-time password systems show promise as a proposed countermeasure^{*79} for the list-based attacks of recent years, through replacement of or combined use with conventional ID/password systems.

■ Differences Between Authentication and Authorization

In the digital world, the ultimate aim is not for a user to login, or for a server to authenticate an entity with a certain ID. After a server identifies a user, authentication is performed to provide appropriate service or enable access to various resources. Once an entity is validated, the act of granting privileges to the entity with the corresponding ID is called authorization, and this is considered separately to authentication. Authorization merely grants the proper access privileges to a given ID. A policy is a prerequisite for this, and there is an approach in which roles are separated into an entity PDP (Policy Decision Point) that grants privileges to a given ID based on the policy, and an entity PEP (Policy Enforcement Point) that actually applies the determined results^{*80}. Access control methods for granting privileges to a given ID include models such as RBAC (Role-Based Access Control)^{*81} and ABAC (Attribute-Based Access Control)^{*82}. These methods take the approach of determining whether to grant privileges based on attribute information associated with an ID, instead of granting privileges to the ID directly. The main difference between these two methods is that RBAC handles attribute information with fixed attribute values called roles, and ABAC adopts a configuration method in which attribute information called an attribute takes a value within a given range, with privileges selected based on the range that this value belongs to. For example, the RBAC approach is to grant access privileges based on position or gender, while the ABAC approach is to grant them based on

*73 NIST Special Publication 800-63-2, Electronic Authentication Guideline (<http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-63-2.pdf>).

*74 The term credentials is sometimes used to refer collectively to all information used in authentication, but for its use in this report we think that clear classification should be made with regard to public information and private information.

*75 ITU-T Recommendation X.509 | ISO/IEC 9594-8, Information Technology - Open Systems Interconnection - The Directory: Public-key and attribute certificate frameworks

*76 Lamport, “Password Authentication with Insecure Communication”. Communications of the ACM 24.11, November 1981, pp.770-772.

*77 N. Haller, “The S/KEY One-Time Password System” (<http://tools.ietf.org/html/rfc1760>).

*78 N. Haller et al., “A One-Time Password System” (<http://tools.ietf.org/html/rfc2289>).

*79 Ministry of Internal Affairs and Communications, “A collection of list-based attack countermeasures” (http://www.soumu.go.jp/main_content/000265404.pdf) (in Japanese).

*80 OASIS, “eXtensible Access Control Markup Language3 (XACML) Version 2.0” (http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-core-spec-os.pdf).

*81 David Ferraiolo, Richard Kuhn, “Role-Based Access Controls”, 15th National Computer Security Conference (<http://csrc.nist.gov/rbac>).

*82 NIST, Attribute Based Access Control (ABAC) Overview (<http://csrc.nist.gov/projects/abac/>).

ages or periods/times. The attribute information used for authorization is sometimes managed by the entity that assigns IDs and performs authentication. However, another method involves an entity that manages attribute information (an attribute authority) issuing credentials and associating IDs and attribute information. The benefit of circulating attribute information as credentials is that it protects against information not required under normal circumstances also being sent to service providers along with attribute information that a user should disclose to gain access privileges. As a result, only the minimal amount of attribute information is disclosed to receive service^{*83}. Figure 19 is a conceptual diagram showing the flow from token-based authentication to the circulation of each type of credential and the granting of access privileges.

■ ID Federation Technology

There are a number of frameworks called ID federation for implementing single sign-on, which enables services in multiple realms to be used at the same time. Here we focus mainly on the SAML (Security Assertion Markup Language)^{*84} approach. The roles in SAML can be divided into (1) the IdP (identity provider) that conducts tasks related to ID assignment and authentication, (2) the RP (relying party) / SP (service provider) that trust and accept the credentials issued by the IdP, and (3) the end user. The end user presents their credentials when using the services of the SP. It is the IdP that issues these credentials. Based on the policy determined by the SP, credentials called assertions are requested, and issued by the IdP. In SAML, the protocol specifies the actions taken when a browser behaves as an end user, in addition to the data format for assertions. Regarding credential (assertion) types, it is possible to circulate information that guarantees attribute information, as well as information related to authorization decisions, in addition to information about authentication results. When doing this, as shown in Figure 19, it is possible to think of the organization issuing credentials relating to authentication as the IdP in a narrow sense, while breaking down the roles into the attribute authority that issues credentials regarding attribute information, and the PDP that issues credentials regarding authorization decisions. Furthermore, the SP serves as the PEP, but depending on the implementation it is also conceivable to omit the issuing of assertions by the PDP and have the SP handle the role of both PDP and PEP.

OpenID^{*85} is another framework based on concepts similar to SAML. The circulation of attributes is enabled using a system in which the RP verifies the credentials (called claims) issued by an entity called an OpenID Provider (OP), which plays the roles of narrowly interpreted IdP and attribute authority mentioned earlier. Operations are possible even when the OP and RP do not coordinate in advance, and an extension specification regarding discovery services for facilitating the discovery of the appropriate OP/RP by each RP/OP is also available. Furthermore, specifications for the delegation of authority have also been established, and OAuth^{*86} provides a system in which users with a token for accessing resources temporarily grant only privileges to a delegate without disclosing the private token information. This may be confusing as the technical terms are different, but users known as resource owners can enable delegates to access resources via a PDP called an authorization server, which issues credentials including authorization information called an access token without passing on private information.

*83 Additionally, there is also the approach of gaining access privileges using a pseudonym for the ID as well, but we do not cover this here.

*84 OASIS, Security Assertion Markup Language (SAML) (<http://xml.coverpages.org/saml.html>).

*85 OpenID Foundation, OpenID specifications (<http://openid.net/developers/specs/>). Previously formulated specification groups such as OpenID Authentication have been discontinued and incorporated in OpenID connect 1.0, which is based on OAuth 2.0.

*86 D. Hardt, "The OAuth 2.0 Authorization Framework" (<http://tools.ietf.org/html/rfc6749>). OAuth 2.0 specifications are established on the following page (<http://tools.ietf.org/wg/oauth/>).

In this way, we expect that a variety of specifications based on new models will continue to be established in the future. We hope that learning about general approaches to IDs, tokens, credentials and authentication, authorization, and access control as presented here will assist in the understanding of new specifications.

1.5 Conclusion

This report has provided a summary of security incidents to which IJ has responded. In this report, we presented countermeasures against the alteration of domain name registration information, and looked at the openioc_scan plug-in that scans for threats lurking in the memory of a device. We also examined ID management technology. IJ 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.

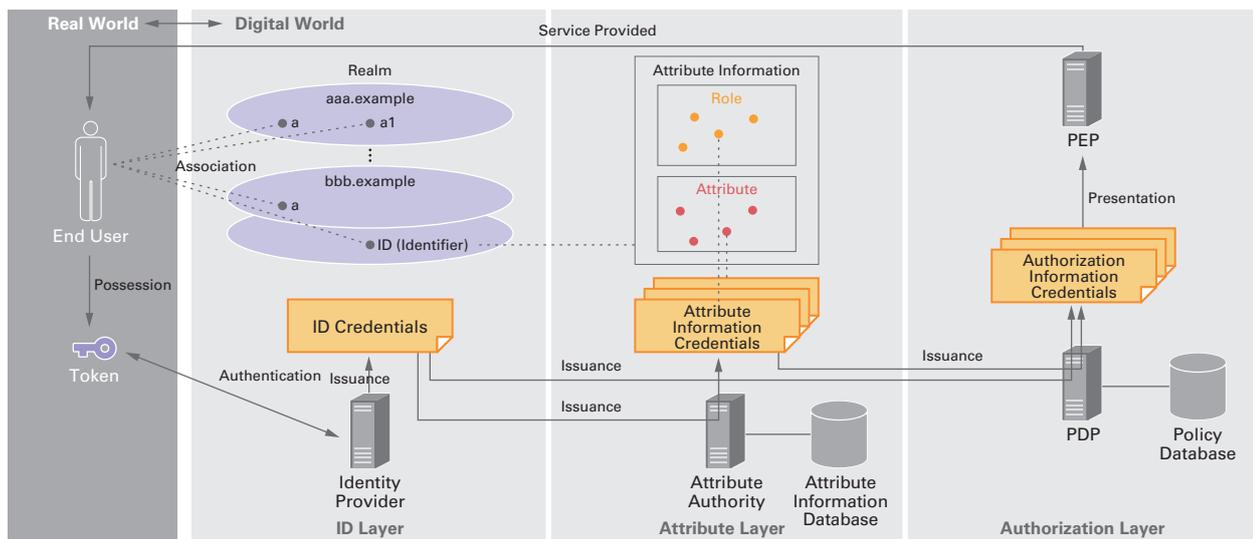


Figure 19: Flow From Token-Based Authentication to Circulation of Credentials and Granting of Access Privileges

Authors:



Mamoru Saito

Manager of the Office of Emergency Response and Clearinghouse for Security Information, Service Operation Division, IJ. After working in security services development for enterprise customers, Mr. Saito became the representative of the IJ Group emergency response team, IJ-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)

Minoru Kobayashi (1.4.1 Countermeasures Against the Alteration of Domain Name Registration Information)

Takahiro Haruyama (1.4.2 The openioc_scan Plug-in That Scans for Threats Lurking in Device Memory)

Yuji Suga (1.4.3 ID Management Technology)

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

Contributors:

Tadashi Kobayashi, Masahiko Kato, Masafumi Negishi, Yasunari Momoi Office of Emergency Response and Clearinghouse for Security Information, Service Operation Division, IJ

Introducing the IJ GIO Cloud Service - Part 2

Last year we presented the first part of our introduction to the IJ GIO Cloud Service.

In this second installment, we examine the status of service utilization in fiscal 2014, and discuss our activities with regard to the development, operation, and maintenance of large-scale infrastructure.

2.1 Introduction

2.1.1 Cloud Market Trends

While the Internet business-oriented usage by parties such as social application providers that has supported the growth and expansion of the cloud market to date is slowing and leveling off, it is evident that enterprise usage at major corporations is growing, and this growth is expected to continue going forward. With these changes in usage trends, more focus is being placed on connecting to (or coordinating with) a variety of cloud platforms as part of a multi-cloud service, rather than a single company providing cloud services. IJ is also paying close attention to these developments, and we are preparing to provide multi-cloud services in addition to our existing multi-carrier connectivity.

2.1.2 Service Overview

■ Topics

We will start offering the IJ GIO Component Service Database Add-on In-memory Platform for SAP HANA® in April 2015. This service enables customers to run a production environment incorporating SAP's SAP HANA® in-memory database in the cloud. We previously launched the IJ GIO for SAP Solution PoC for SAP HANA® (Figure 1) cloud environment for prior validation in June 2014, and this new service provides the same features for production environments. IJ has obtained SAP

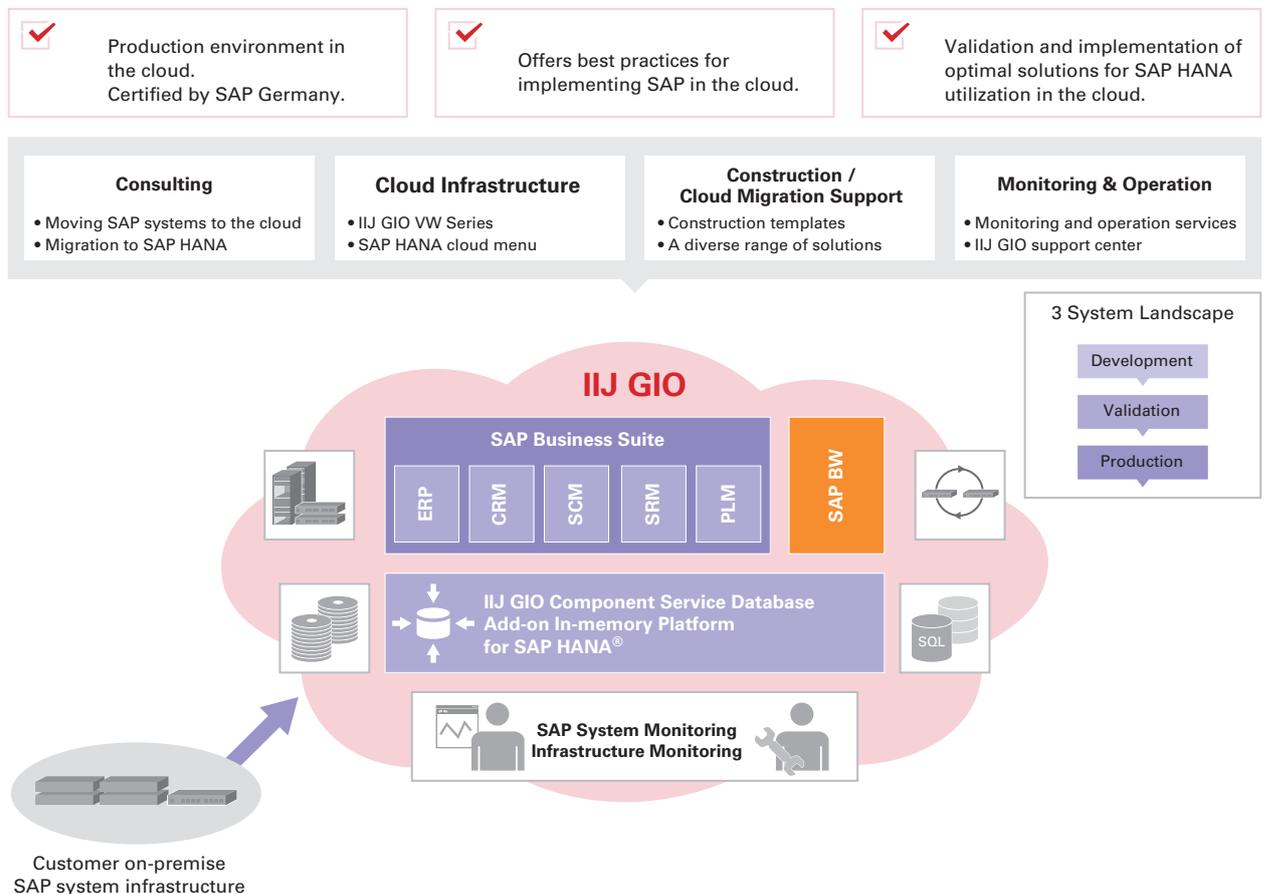


Figure 1: IJ GIO for SAP Solution

Certified in Hosting Services and SAP Certified in Cloud Services*1 certification, and provides worry-free access to a high level of services and security in SAP production environments. This keeps initial investment and operation costs down, and enables swift development and deployment of systems to quickly respond to business decisions.

In January 2015, we began offering the IJ Cloud Exchange Service for Microsoft Azure (Figure 2). This uses the Microsoft Azure ExpressRoute walled garden network service, which IJ is the first partner in Japan to offer. Until now connections to Microsoft Azure were made only over the Internet, but ExpressRoute enables connection via a walled garden network. Starting with this service, IJ plans to further expand its multi-cloud services.

■ Service Facility Development Status and Usage Trends

Regarding the status of facility development for IJ GIO in fiscal 2014 (as of March 2015), we forecast around a 20 percent increase in physical server numbers and a 40 percent increase in physical storage space, compared to the end of fiscal 2013. As far as usage trends are concerned, there continues to be fast-moving fluctuations each month in the use of virtual servers used mainly as front-end Web/AP servers for social application providers. Meanwhile, regarding physical servers chiefly used as database servers, we have seen further progress in the consolidation and shift from servers using basic local hard drives or FC storage to physical servers equipped with semiconductor storage that offer high I/O performance. In the area of enterprise-oriented equipment, this was a year that saw remarkable growth in storage over and above server numbers. The dramatic growth in storage usage suggests that major corporations requiring large amounts of storage are migrating mission-critical systems to the cloud, and real utilization of systems migrated to the cloud is progressing. We expect these trends to continue in fiscal 2015.

2.2 Large-Scale Infrastructure Technology Initiatives

2.2.1 Infrastructure Development

■ Cloud Infrastructure Considerations

As a cloud service provider, we are always thinking of ways to operate large numbers of hardware devices efficiently. If we were simply focused on optimization, then the homogenization of hardware would be ideal, but in reality this is still difficult

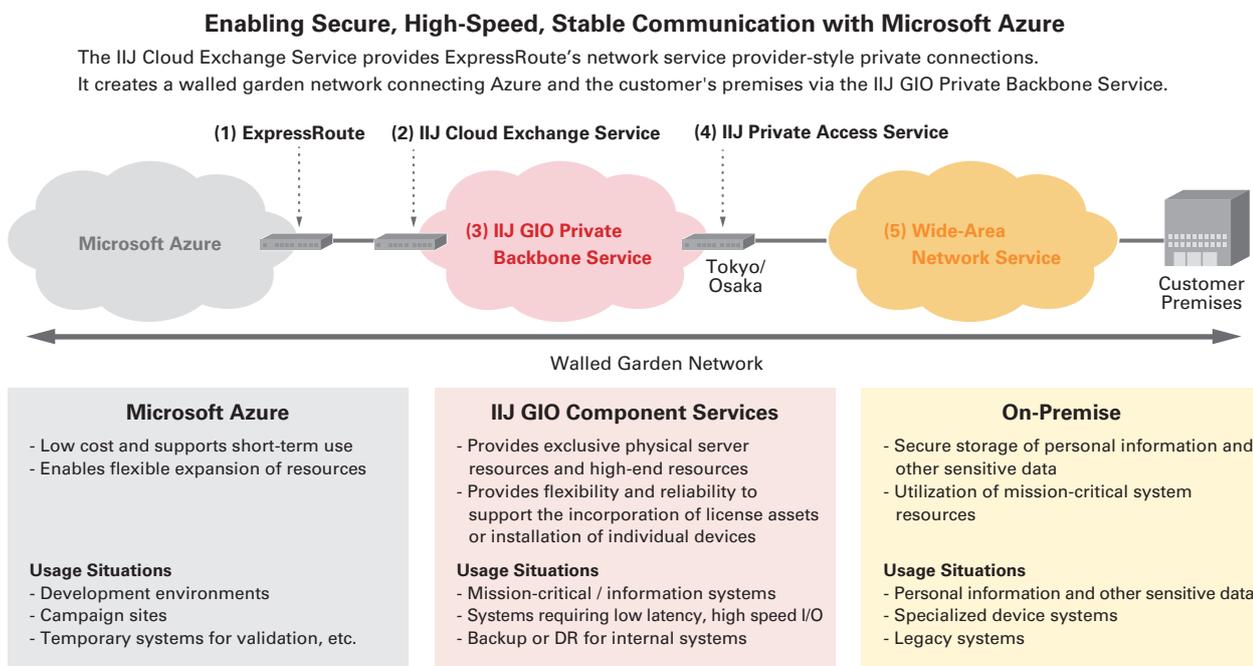


Figure 2: IJ Cloud Exchange Service for Microsoft Azure

*1 Meets the service quality criteria established by SAP Germany, and is recognized as a SAP-certified cloud service provider and SAP-certified hosting service provider (certified April 18, 2013).

at present considering IJ GIO's wide-ranging lineup of services, as well as its cost-effectiveness. As a result, we are currently aiming to implement only as much homogenization as possible.

For example, infrastructure for the public cloud-oriented IJ GIO Hosting Package Service targeted at Web-based companies and social application providers must compete in the public cloud field where price competition is fierce, so we provide practical-minded hardware configuration and OSS combinations that eliminate excess redundancy. Meanwhile, infrastructure for the private cloud-oriented IJ GIO Component Services targeted at the enterprise sector incorporates hardware redundancy at least equivalent to on-premise solutions, as on-premise expansion resources or a receptacle for migration from on-premise equipment. We also provide combinations incorporating commercial products to meet requirements for cloud infrastructure that guarantees operation with existing commercial products that a customer coordinates with or brings in. Above all else, it is essential to diversify technological risks to provide stable service infrastructure. Because it is not possible to completely eliminate the risk of issues with firmware or drivers, it is best to avoid using only products from a specific vendor even for the same component, to localize the impact to a certain extent should an issue occur. In other words, multi-vendor systems have benefits from the perspective of diversifying technological risks. Adopting multiple vendors also enables you to select from a range of vendor products, making choices based on commodity technology instead of being dependent on specific vendor product technology. As a result, the principle of competition applies when procuring equipment, contributing to reduced procurement costs. Adopting commodity technology also leads to optimizations in operational areas such as delivery and maintenance, making integrated management using common management tools possible.

In addition to diversifying technological risks, we are pushing forward with standardizing the base equipment for different service infrastructure with an eye toward homogenization. This makes it easy to reuse equipment in different service infrastructure by simply adding or changing components. In our current service development, this standardized approach is applied to the physical design of service infrastructure as long as there are no special requirements. Standardization enables the surplus inventory for each service infrastructure to be treated as a shared inventory pool, also contributing to reduced inventory risk. With regard to the hardware adopted, we have seen an increase in the range of options available centered around overseas cloud service providers over the past few years. In the server field, in addition to the purchase of general server vendor products as before, there have also been cases in which requests for manufacturing were made to ODM vendors, with a central focus on Taiwan. Servers produced through knowledge gained from the Open Compute Project (OCP) and certified by it have started appearing on the market as well.

Except for a few services, IJ mainly uses general-purpose IA servers supplied by server vendors for the IJ GIO cloud service infrastructure developed in-house. Of course, we have also evaluated and validated server products from ODM vendors. However, as a result of overall comparison of our service infrastructure requirements and the economies of scale gained from procurement, we came to the conclusion that server vendor products are more suitable for our current cloud service infrastructure. That said, this may not always be optimal in the future, so we will continue to keep tabs on trends in the aforementioned server vendor products, ODM vendor products, and OCP servers to respond to changes in the market swiftly.

With regard to software, we believe it is inefficient to create it all from scratch ourselves, and this would have few benefits for users. From our perspective, there is value in being able to provide services by adopting good products on the market based on customer needs, whether commercial products or open source software (OSS). However, there are also many difficulties associated with market products, such as the need to refine them when they don't support the scale required by cloud service providers. There is a lot of talk in public circles about how to handle Windows Server 2003, for which extended support will end on July 15, 2015, but the end of support is of course always an issue with commercial products, and even OSS. In light of this issue, for IJ GIO we share information with and issue alerts to users of services that incorporate products for which support is scheduled to end. In particular, the private cloud-oriented IJ GIO Component Services targeted at the enterprise segment provide a range of services that combine a large number of commercial products, and support lifecycles are defined for each vendor. To ensure that operation of these is always guaranteed and support can be provided, we list the support lifecycles of products for each service in the form of a service lifecycle, and periodically update information to help users continue using services safely and securely. The needs of users are not met by merely providing services at low cost, and IJ takes this into account.

■ The Accumulation of Ongoing Technological Development

Infrastructure development is triggered by the dual action of service development requirements and accumulated technological development. In particular, developing the optimal infrastructure to meet service development requirements hinges upon how much day-to-day knowledge and experience you have in a variety of technological development fields (Figure 3). As an organization, IIJ continually repeats a technological survey cycle that involves gathering information on product trends and roadmaps for each manufacturer, comparing and evaluating technology and trends, and producing periodic reports. We also refer to feedback from the development site and management based on these reports, and perform internal product validation to assess whether products that show promise would stand up to service operation (Figure 4).

Our service development team plans and develops services in keeping with the technical knowledge we have accumulated and our business strategy. In the design phase infrastructure is designed based on the validation results produced by the product validation team, with the aim of guaranteeing service quality and reducing design time. Surveys are of course targeted at servers themselves, but also cover areas as diverse as hardware such as storage and network equipment, software such as OSES, hypervisors and SDNs, and control systems such as OpenStack and CloudStack.

Examples of our ongoing initiatives include keeping pace with server-side flash and Intel’s next-generation CPUs. Each time a new product is released, IIJ obtains evaluation units from the vendor and works on identifying changes in performance characteristics over generations. We also clarify the impact on service and points of concern stemming from switching to successor models, and develop internal guidelines to facilitate their smooth introduction to the service environment. By continually tracking the above technological points even for comparatively new products and models with a short product

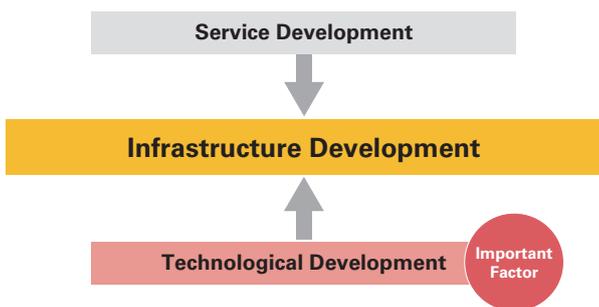


Figure 3: Infrastructure Development Approach

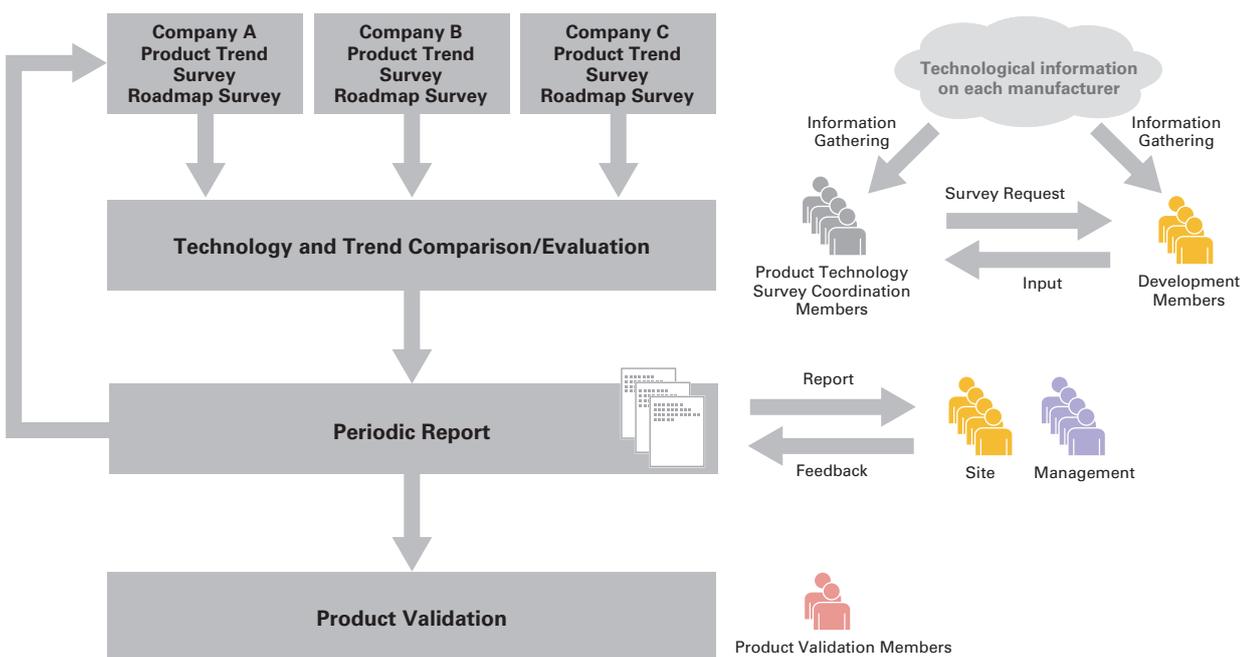


Figure 4: Technology Trend Survey Process

cycle, we can rapidly incorporate them into our infrastructure in response to service requirements and provide them. One of our initiatives in the area of software is participation in Microsoft's Technology Adoption Program (TAP). By participating in TAP, we can cultivate a better understanding of the design objectives, new features, and changes from the current version for the next versions of the Windows Server operating system and System Center management products. In addition, we can further our knowledge of implementing and upgrading these products. IJ also submits feedback in the form of feature requests and notes on specification improvements to meet our service development requirements. Through these efforts, we are able to add new OS and hypervisor versions to our infrastructure and provide services immediately after their release. Having our requests as a cloud service provider incorporated makes it possible to migrate or run services with minimal downtime each time a new version is released.

2.2.2 Infrastructure Operation

Infrastructure operation is just as important as infrastructure and technology development. Here we discuss infrastructure maintenance activities such as failure and fault handling that are the mainstay of infrastructure operation. As with on-premise equipment, system failures also occur on the equipment of cloud service providers. However, even though we are not using new products straight after their release, extremely rare system failures with no precedent may be encountered due to the nature of multi-tenant environments, which feature large numbers of users applying a range of configurations and uses. Instead of simply dealing with rare failures on the spot, we check whether the failure affects the infrastructure as a whole, and monitor the situation for a set period to guard against recurrence. Although identifying the cause of a failure requires money and time, inaction during the initial response can lead to large-scale failure, so our stance is to aim for establishing surefire countermeasures.

■ Infrastructure Maintenance Activities Such as Failure and Fault Handling

Examples of infrastructure maintenance activities include proactive maintenance, measures for dealing with underlying causes, and replacement.

Proactive maintenance includes preventive maintenance for potential and fatal faults discovered after a service is released, as well as the application of patches to keep pace with the support lifecycle of firmware or software. It is implemented as part of the maintenance for receiving appropriate vendor support in case of unanticipated events. For information gathering, general security information is issued from an internal information sharing site daily, and we have a framework in place for checking failure and vulnerability information for the products and technology used in service infrastructure on an ongoing basis. The same applies to support lifecycles. However, we don't merely take vendor information on faith. We first carefully validate whether there is any impact on the existing environment, including parts we have built ourselves, using validation/staging environments practically identical to the production service environment. Then we select the parts we consider important for the stable operation of the service environment. It is also rare for the systems or tools provided by vendors to take deployment in large-scale infrastructure into consideration. In many cases, firmware and patches are provided on the assumption that they will be applied individually to equipment step-by-step, so it is necessary to devise a method such as automation to apply them reliably to large quantities of equipment. This demonstrates that proactive maintenance strongly resembles the technological development process detailed above. Dealing with underlying causes involves taking measures to tackle failures occurring after the release of a service that impact the infrastructure as a whole. As mentioned earlier, we sometimes encounter extremely rare system failures without precedent. Even for unprecedented failures, it is necessary to recreate the failure to resolve it. There is no way for vendor support to take measures without recreating the failure, leading to a dead end. However, due to vendor support performing reproduction tests using only a handful of validation units, in some cases a failure cannot be recreated even after we convey the details of the phenomenon. This is because some failures have an extremely low chance of occurring, so reproducing them requires many repeated attempts. Consequently, we occasionally have vendor support engineers visit us to confirm the failures actually occurring, and have ongoing reproduction tests carried out. We are able to reproduce even failures that normally only reoccur after a hundred attempts on ten machines comparatively easily because we can execute a range of patterns at once utilizing large amounts of equipment. Reproduction tests are performed by IJ engineers who have accumulated reproduction knowledge through encountering many failures in past operational experience. Taking advantage of this superior reproduction capability, we constantly endeavor to encourage vendor support to have fixes made promptly to create more stable service infrastructure. For example, when introducing a large quantity of server equipment in the past, there were frequent NIC communication failure incidents right after initial installation. Because vendor support insisted that there had been no other cases in the past, we analyzed the failure events, established steps to reproduce them, and provided vendor support with information on environments the failures occurred in as well as the steps taken. This enabled vendor support to confirm the failure events,

and a few months later the cause was identified. Vendor support came up with every workaround for the failures they could think of, such as updating the BIOS, firmware and drivers, as well as changing the settings, but ultimately the failures did not stop occurring. In the end, the vendor determined it would be difficult to avoid the failures with the existing NIC, so the radical measure of switching to a different manufacturer's NIC was taken.

We work on discovering and eliminating the risk of widespread failures on a daily basis, based on the reproducibility of faults like this. Due to an accumulation of frustrating experiences in infrastructure operation such as the one above, the diversification of technological risks is an indispensable requirement for service infrastructure today.

Replacement involves the renewal of equipment when devices you are using reach the end of their support lifecycle. For corporate IT systems, hardware is generally replaced in cycles of three to five years when the maintenance or lease expires. Service providers like us that offer cloud services also procure equipment (with some exceptions) from general IT hardware vendors, and replacement will be necessary after a set time (although this can be extended a little) based on the product lifecycle determined by each hardware vendor when equipment is installed. For example, the IIJ GIO Hosting Package Service infrastructure that provides Linux-based virtual servers faced the problem of maintenance support expiring for the storage devices it used in expansion disks. The existing storage devices required renewal due to end of manufacturer support and lease expiration. Cloud service infrastructure always calls for zero downtime when replacing equipment, but the existing models for these storage devices presented a catch-22 situation, as they did not support storage migration to different models, and the new models had no migration function to begin with. To solve this problem we collaborated with the hardware vendor. By meeting with the vendor many times over a number of months, we convinced them to add a function to existing models that enabled migration to different models. Going one step further, we also gained official support for this migration function by having the function added through the modification of publicly available firmware, instead of asking for the migration function to be provided to us individually. However, the problem was not solved by this. Another issue that came to light was that migration work for each chassis required the execution of several hundred commands. Overall, several tens of thousands of commands would need to be executed. Because there was no way to execute these manually in an error-free manner, we prepared a system for executing the migration function reliably by creating a new semiautomated script.

As the case above shows, collaborating with hardware vendors and taking user requirements into account are areas that require experience and knowledge. These kinds of responses, including the semiautomation of processes, are likely to be missed and thus fail to draw much attention from users, but they are hard to obtain for individual customers. I believe that aspects such as this are one of the benefits of using a cloud service provider like us.

2.3 Conclusion

In this report we have given an overview of the infrastructure technology that supports the IIJ GIO service. Hopefully this gives customers that use this service a glimpse of what goes on behind the scenes. We will continue to strive for day-to-day stable operation and improved quality in the IIJ GIO service, to provide safe and secure cloud services to everyone.

Author:



Shinri Kimura

Manager, System Infrastructure Technology Department, Platform Division, IIJ. As service manager of the IBPS service that preceded the IIJ GIO cloud service, Mr. Kimura was responsible for its development and operation. He later participated in IIJ GIO from the planning and concept stage, and he is now engaged in the development and operation of IIJ GIO system infrastructure services.

Contributor:

Tsuyoshi Inaba, System Infrastructure Technology Department, Platform Division, IIJ.

UKAI Virtual Disk Storage

This article introduces our research into the UKAI*¹ storage system, which enables flexible control over the actual data locations of virtual disk images of virtual machines for which location management is difficult.

3.1 Infrastructure with Virtualization as a Prerequisite

The word virtualization no longer has a fresh ring very often these days. Virtualization technology now occupies an important position as a component technology for supporting services. Of course, virtualization technology itself is not new at all, having been used in a range of situations since computers first came into practical use. Even the multitasking technology implemented in current mobile phone OSes was first conceived for the parallel use of expensive physical CPUs among multiple programs via virtualization. The Java language widely used in enterprise services could also be called a type of virtualization technology that utilizes Java virtual machines. Despite this, virtualization has attracted a great deal of attention in recent years, because the prospect of executing the server environments we use every day as virtual machines at an acceptable speed is within sight.

Some network-based services have already been composed of virtual equipment thanks to the maturity of the operation technology of virtual machines. The use of virtualization technology allows for the swift deployment of services, while also making it comparatively simple to expand resources urgently to cope with high loads, and reduce resources when usage drops off, which would be difficult if physical machines were used. Performance is of course inferior to using physical machines directly, but we are now at a point where this gap is evened out by the ease of operation. Virtual machines will never catch up to the performance of physical machines, but based on advances in basic performance, it is not hard to imagine that virtual environments will one day meet the level of quality required by many services. In the same way that applications that monopolize CPU resources have all but disappeared, it is likely that installing service OSes directly on a single piece of hardware will one day be a thing of the past. We will soon be in an age where virtual machines are a fundamental component in service infrastructure, other than a few special applications.

The first step in operating virtual infrastructure efficiently is to construct data centers dedicated to virtual machines. IJ's Matsue Data Center Park*² is one of these data centers, and this is used as the infrastructure for the IJ GIO cloud service. One way to raise the efficiency of a data center is to increase its scale. However, in a country with limited space such as Japan, it

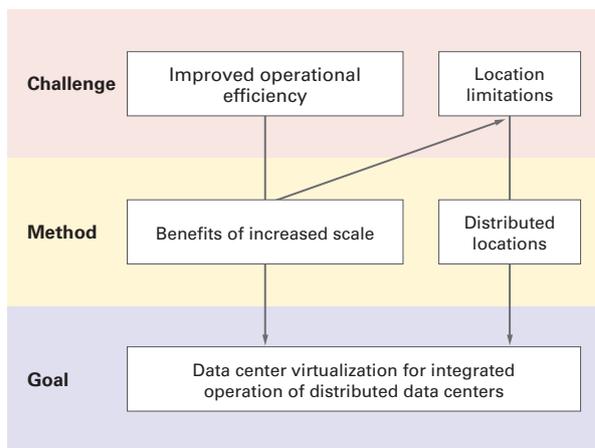


Figure 1: Virtual Data Center Demand and Challenges

*1 Keiichi Shima, "UKAI: Centrally Controllable Distributed Local Storage for Virtual Machine Disk Images", In Proceedings IEEE Globecom 2013 Workshop on Cloud Computing Systems, Networks, and Applications (CCSNA), December 9, 2013.

*2 IJ's Matsue Data Center Park (<http://www.ij.ad.jp/DC/dcpark/>) (in Japanese).

is hard to make a significant investment and construct massive data centers when commensurate demand is not immediately anticipated. This leads to the concept of virtual data centers, in which geographically separate medium-scale data centers are operated as a virtual large-scale data center. When configuring infrastructure services across distributed data centers, virtual resources must be arranged and operated efficiently, taking into consideration geographical conditions and network latency (Figure 1). Systems for the flexible location and relocation of virtual resources will be increasingly important in the future.

3.2 The Necessity of Disk Storage for Virtual Machines

Broadly speaking, virtual machines are comprised of three resources. The first is the CPU and memory that form the core of a computer, the second is the network for interconnection between computers, and the last is storage for retaining systems and data. Technology that enables the flexible location and relocation of these resources will contribute greatly to the operation of adaptable virtual infrastructure in the future.

As you are no doubt aware, relocation of the first resource (CPU and memory) is already possible at a practical level. The virtual hardware environment that is a prerequisite for virtual machines can be unified to a certain extent, so as long as virtual machine specifications are decided, the hypervisor that executes them can be operated anywhere. Some virtualization technologies also provide a live migration function that migrates virtual machines to another hypervisor while they are running.

The research and development of technology for relocating the second resource (networks) has been underway for some time, with Software Defined Networking (SDN) technology being one of the strongest candidates. VLAN was a form of network virtualization used as standard up until now, but due to the increase in the scale of data centers, it is nearing the bounds of its capacity. Use of SDN technology enables networks to be sectioned out in virtual machine units, allowing for flexible network configurations that were difficult to implement using VLAN in the past.

This leaves the last resource, storage, but at this point in time there are no good options available. NFS and iSCSI are often used as technologies for providing virtual disks to virtual machines. Considering factors such as the granularity of disk volume management and redundancy of the system in case of disk failures, in most cases it is likely that iSCSI products are selected for large-scale infrastructure. This storage is connected via a network, so as long as the network resources of a virtual machine are correctly relocated, it should still be possible to access its storage in theory. However, the data kept in this storage remains fixed to a physical storage server. Considering the prospect of distributed operation based on the concept of virtual data centers in the future, it is possible that when a suspended virtual machine is restarted, it may restart at a different data center depending on the status of hypervisor resource allocation. When storage resources are allocated to a fixed device somewhere, it affects the performance of virtual machines launched in remote locations. We need to be able to flexibly relocate storage resources in the same way as CPU and network resources, but NFS and iSCSI cannot fulfill this requirement.

At the IJ Research Laboratory, we are working on the research and development of the UKAI storage system that enables flexible control over the actual data location of virtual disks used with virtual machines, taking into consideration environments for the integrated operation of distributed data centers as a virtual data center.

3.3 UKAI: Location Aware Virtual Disk Storage

UKAI is designed to achieve the following three goals.

1. Administrators are able to control the location of actual data for virtual disks based on operational decisions
2. Functions practically in environments where network latency is not uniform
3. Has redundancy in the event of failures

The actual data of virtual disks is stored in storage nodes distributed across various data centers. When creating a virtual machine, the actual data of the virtual disk allocated to that virtual machine will ideally be located close to the place where that virtual machine is running, whenever possible. It is necessary to reduce latency between the virtual machine itself and the actual data of the virtual disk, and prevent performance degradation in a distributed environment, by situating storage nodes at each data center and specifying the location of the actual data for a virtual disk operationally. However, as mentioned earlier, virtual machines are not always located on the same hypervisor on a permanent basis. Virtual infrastructure is accessed and shared by multiple users. Depending on service usage, there may be cases in which you have no choice but to launch a virtual machine on another hypervisor, or even a hypervisor at another data center at times. The minimum requirement is being able to provide access to virtual disks transparently even in these situations. However, accessing storage resources at another data center dramatically degrades virtual machine performance. UKAI separates the management of virtual disks and actual data, enabling the actual data to be migrated while the virtual disk information remains fixed. This makes it possible to freely change the location of the actual data comprising a virtual disk, while the virtual machine is running, based on operational decisions.

Dealing with network latency is an important issue for distributed environments. It is particularly crucial to estimate the impact of latency when providing services across locations distributed over a wide area, as is the case with virtual data centers. In general, there are two types of latency when operating over a wide area. The first is communication latency within a data center. This applies when multiple storage nodes are in operation, and these nodes are located within the same data center. In this case, latency is minimal (around 1 ms or less), and we can expect it to be uniform. The other type is communication latency between data centers. This varies depending on the relative positions of the data centers, and we can expect latency to have greater fluctuations compared to communications within a data center. UKAI leaves the selection of the location for actual data up to the administrator, and operation is carried out after predicting the latency that will occur. Even when launching a virtual machine in a remote location, the degree of performance degradation can be predicted as long as the latency between that virtual disk and the actual data is quantitatively known. A migration plan is then formulated for the actual data based on this.

The last goal of redundancy goes without saying. You cannot afford to have the whole service shut down due to a node failure somewhere when carrying out the distributed operation of storage nodes. The same applies to situations in which access to service nodes is lost due to network failure. It is of course impossible to protect against all failures, so the best approach is to anticipate a certain number of failures, and add redundancy to cope with them. For UKAI we anticipate disk and network hardware failures on storage nodes. With regard to device failures within data centers (switch failures, etc.), we assume that redundancy for data center equipment is implemented separately. Because virtual disks and actual data are managed separately in UKAI, it is possible for a certain virtual disk to have multiple copies of a piece of actual data. Placing multiple copies within the same data center enables them to be run as network-based mirror disks. Storing copies in other locations also makes it possible to operate them as a disaster recovery measure, although storage access performance would drop.

3.4 Implementing UKAI

We have released a prototype implementation for validation of the UKAI design^{*3}. This implementation takes into consideration the following points.

- Compatibility with existing hypervisors
- Elimination of single points of failure through distributed systems
- Interfacing with cloud controllers

There are two conceivable methods for providing virtual machines with virtual disks. The direct method involves extended implementation of a new virtual disk format on the hypervisor itself. Directly embedding this is likely to lead to better performance, but it is necessary to develop separate implementations for each hypervisor. Another method involves providing virtual disk images as files. The simplest method many hypervisors have is a system in which a file is treated as a single virtual disk image and provided to the virtual machine. UKAI adopts this latter method, constructing virtual disk storage independent of the types of hypervisors.

When designing a distributed system, the greatest concern is how much redundancy should be implemented for when failures occur. As mentioned earlier, UKAI is able to operate multiple storage nodes, so storage node failures can be coped with to a point. Another issue is how to save the information that constitutes a virtual disk (virtual disk metadata) as a distributed system. In our prototype implementation, the Apache ZooKeeper system for distributed coordination support is used.

Figure 2 shows the module configuration of the prototype system. The hypervisor uses files created on a file system that is mounted on a local system as virtual disks. From the hypervisor's perspective, virtual disk images are seen merely as files, so

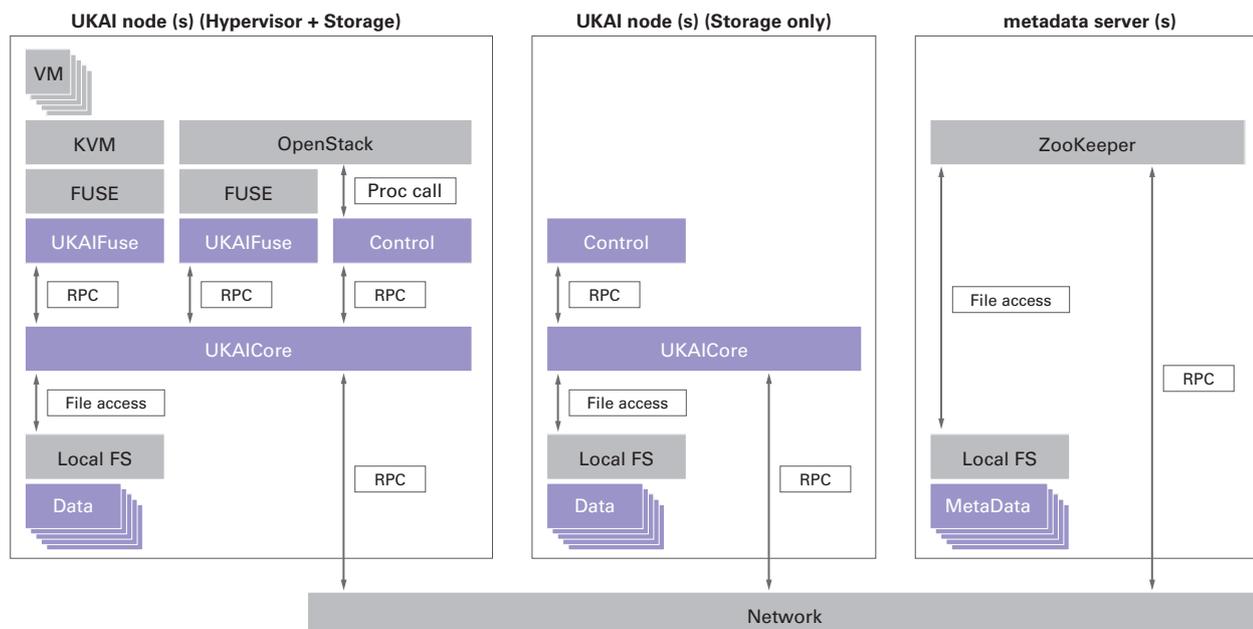


Figure 2: UKAI Prototype System Module Configuration

*3 UKAI: A Location-Aware Distributed Storage Software for Virtual Machine Disk Images (<https://github.com/keiichishima/ukai>).

they can be operated in the same way as when using local files as virtual disks. UKAI is actually present at the mount point, so file input-output is all intercepted by the UKAI subsystem. In the prototype, FUSE*⁴ and the fusepy*⁵ Python binding for it are used to implement the UKAI subsystem. UKAI itself is also implemented using Python.

As their name suggests, virtual disks don't actually exist, and are merely a collection of pointers to the actual data. Figure 3 shows the relationship between virtual disks and the actual data and storage nodes. As this figure indicates, virtual disks are segmented into multiple data blocks, with each data block having pointers to the actual data. Each data block can have pointers to multiple pieces of actual data, to cope with failures on storage nodes.

The relocation of virtual disks is carried out as shown in Figure 4. Let us assume that when a virtual machine is migrated to or restarted on a different hypervisor, a large amount of communication latency is anticipated between it and the storage node containing the actual data. In this case, a copy of the actual data is first created on a storage node close to the hypervisor to

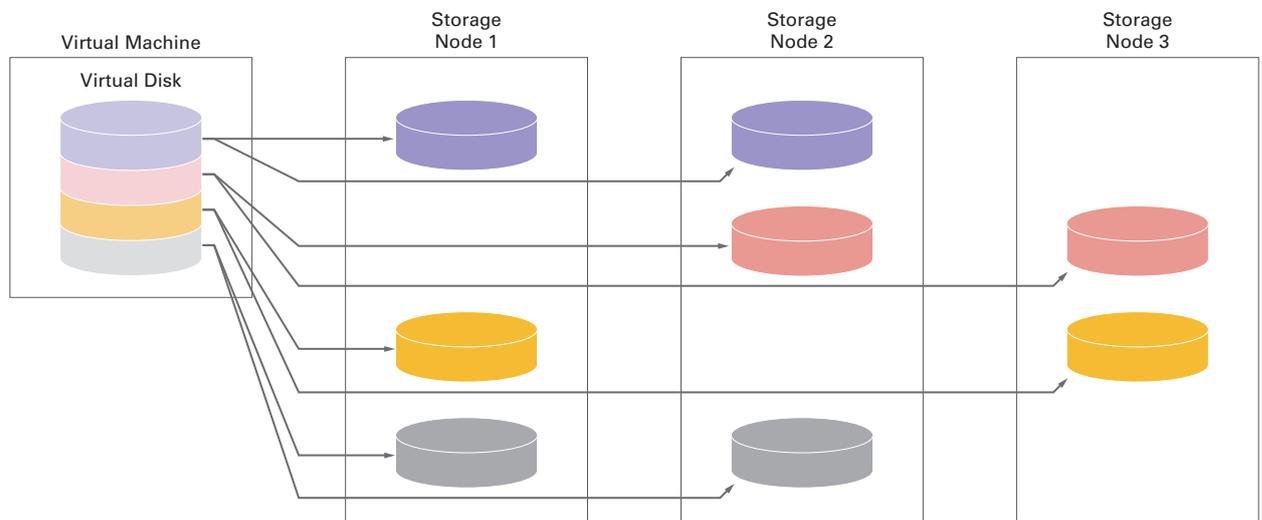


Figure 3: Conceptual Diagram of UKAI Virtual Disk and Actual Data Storage

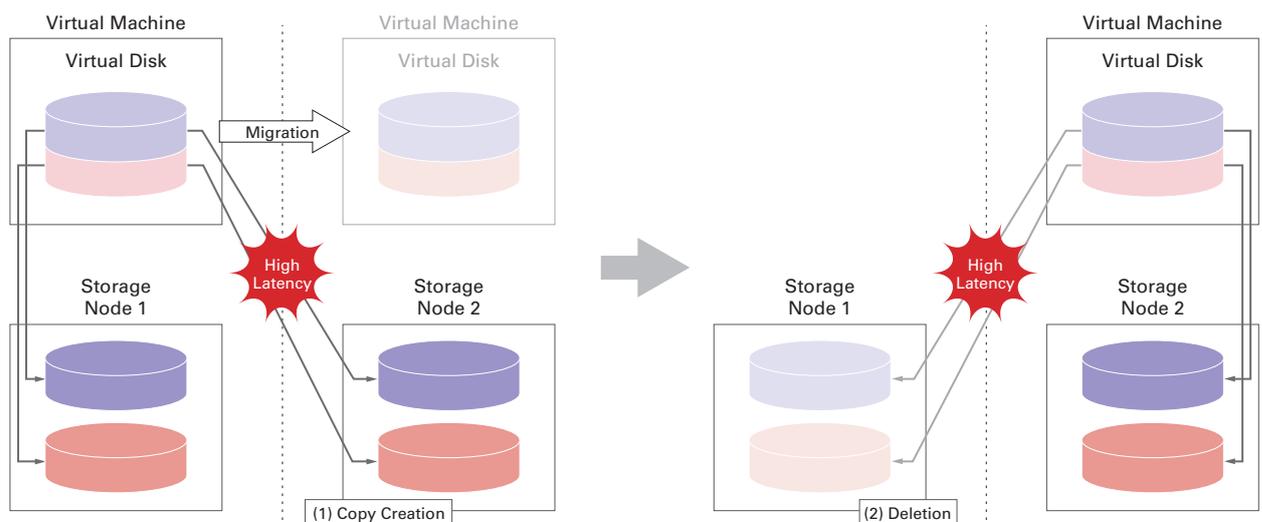


Figure 4: Virtual Disk Relocation

*4 FUSE (<http://fuse.sourceforge.net>).

*5 fusepy (<https://github.com/terencehones/fusepy/>).

which the virtual machine is being migrated. After the virtual machine is migrated, the actual data on the storage node with the higher latency is deleted. Because virtual disks are simply collections of pointers, there is no need for a virtual machine in operation to be aware that a copy of the actual data was created, or that the actual data further away was deleted. Although disk access performance is degraded until relocation is complete, it is possible to optimize the location of the actual data for a virtual disk without stopping the virtual machine.

3.5 Conclusion

As virtual environment performance increases, we believe that many physical machines will be replaced by virtual machines. This will of course have a great impact on services that have continued since the age of physical machines, including server rental, but its true value will be demonstrated when virtual environments that enable the software-based addition or deletion of machines come to be used as components in service delivery models such as SaaS and PaaS. With service infrastructure moving into the cloud, technology for the flexible location and operation of component resources will be essential. Here we have introduced our research on a technology for the relocation of storage, a resource for which location management is difficult. IJ will continue to pursue technological innovations that serve as infrastructure for stable Internet services.



Keiichi Shima

Senior Researcher, Research Laboratory, IJ Innovation Institute Inc. Mr. Shima is engaged in the research and development of virtual machine architecture in wide-area distributed computing environments, and technology for the flexible location of virtual resources.

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.

©2008-2015 Internet Initiative Japan Inc. All rights reserved.

IIJ-MKTG020XA-1503CP-00001PR