Preparing for DDoS Attacks

In this report, we will explain incidents that occurred between July and September 2010, and also examine preparations to be made for DDoS attacks on small-scale systems, discuss security considerations for shared systems such as those on cloud computing environments, and give an overview of digital forensics.

1.1 Introduction

This report summarizes incidents to which IIJ responded, based on general information obtained by IIJ 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 IIJ has cooperative relationships. This volume covers the period of time from July 1 through September 30, 2010. In this period a number of vulnerabilities related to Web browsers and their plug-ins continued to be exploited. There were also incidents that led to financial damages in which SIP was exploited to make toll calls, and in September a series of synchronized attacks on multiple Web servers in Japan occurred in response to the social situation. As seen above, the Internet continues to experience many security-related incidents.

1.2 Incident Summary

Here, we discuss the IIJ handling and response to incidents that occurred between July 1 and September 30, 2010. Figure 1 shows the distribution of incidents handled during this period*1.

■ Vulnerabilities

During the period a large number of vulnerabilities were discovered and fixed in Microsoft Windows*2*3*4*5 as well as applications such as Adobe Systems' Adobe Reader and Acrobat*6*7, Adobe Flash Player*8*9, and Apple's QuickTime*10. Several of these vulnerabilities were exploited before patches were released. A vulnerability in the Linux kernel*11 was also patched. Fixes were also made to vulnerabilities in server applications such as BIND DNS servers*12, and ISC DHCP servers*13, in addition to a number of vulnerabilities in router products such as Cisco Systems' Cisco IOS*14*15. Vulnerabilities were also patched in Apple's iOS*16, which is used as firmware for devices such as mobile phones.

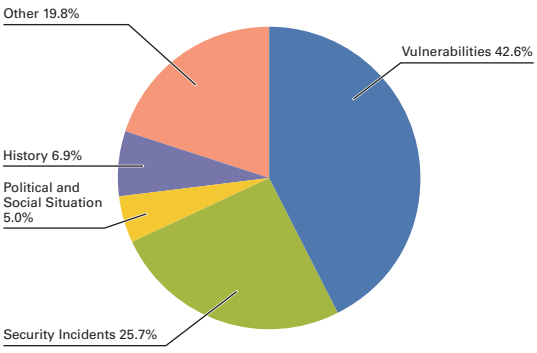

Figure 1: Incident Ratio by Category (July 1 to September 30, 2010)

■ Political and Social Situations

IIJ pays close attention to various political and social situations related to international affairs and current

*1 Incidents discussed in this report are categorized as vulnerabilities, political and social situation, history, security incident and other. Vulnerabilities: Responses to vulnerabilities associated with network equipment, server equipment or software used over the Internet, or used commonly 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: Wide propagation of network worms and other malware; DDoS attacks against certain websites. Unexpected incidents and related response. Other: Security-related information, and incidents not directly associated with security problems, including highly concentrated traffic associated with a notable event.

*2 Microsoft Security Bulletin MS10-042 - Critical: Vulnerability in Help and Support Center Could Allow Remote Code Execution (2229593) (http://www.microsoft.com/technet/security/bulletin/ms10-042.mspx).

*3 Microsoft Security Bulletin MS10-046 - Critical: Vulnerability in Windows Shell Could Allow Remote Code Execution (2286198) (http://www.microsoft.com/technet/security/bulletin/ms10-046.mspx).

events. During this period we turned our attention to the incident in which a Chinese boat collided with Japan Coast Guard patrol vessels in early September.

History

The period in question included several historically significant days on which incidents such as DDoS attacks and website alterations have occurred. During the period for the current report there were warnings of an attack on September 18 (the date of the Manchurian Incident), and we paid close attention to our equipment and our customer networks for attack behavior. Attacks linked to this incident included DDoS attacks and alteration attempts against the websites of a number of organizations related to government agencies, general companies, and unrelated associations.

Security Incidents

Unanticipated security incidents not related to political or social situations were discovered in the form of malware*17 that targets Siemens' industrial control systems*18. An increase was also observed in the unauthorized SIP communications*19 that have been occurring of late. Additionally, a cross-site scripting vulnerability in Twitter was discovered*20 and exploited*21, and a vulnerability in ad distribution servers was exploited*22 to partially alter data and induce users to download scareware.

Other

Regarding trends for other security-related topics, signature protection for the DNS root zone was put into effect*23 to facilitate the implementation of DNSSEC, and in Japan it was announced that DNSSEC will be implemented for JP domain name services in January 2011*24. Microsoft released a patch*25 implementing RFC5746, which was established due to a vulnerability in the TLS protocol relating to the renegotiation feature. In September there was an activity that set out to release information about unpatched vulnerabilities each day, and this led to a large amount of vulnerability information being made public*26.

*4 Microsoft Security Advisory (2269637) Insecure Library Loading Could Allow Remote Code Execution (http://www.microsoft.com/technet/security/advisory/2269637.mspx).

*5 Microsoft Security Bulletin MS10-070 - Important: Vulnerability in ASP.NET Could Allow Information Disclosure (2418042) (http://www.microsoft.com/technet/security/bulletin/ms10-070.mspx).

*6 APSB10-17 Security updates available for Adobe Reader and Acrobat (http://www.adobe.com/support/security/bulletins/apsb10-17.html).

*7 APSB10-21 Security updates available for Adobe Reader and Acrobat (http://www.adobe.com/support/security/bulletins/apsb10-21.html).

*8 APSB10-16 Security update available for Adobe Flash Player (http://www.adobe.com/support/security/bulletins/apsb10-16.html).

*9 APSB10-22 Security update available for Adobe Flash Player (http://www.adobe.com/support/security/bulletins/apsb10-22.html).

*10 About the security content of QuickTime 7.6.7 (http://support.apple.com/kb/HT4290).

*11 A vulnerability was found in the Linux 64-bit kernel. This information is managed as CVE-2010-3081 (http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2010-3081).

*12 RRSIG query handling bug in BIND 9.7.1 (http://www.isc.org/software/bind/advisories/cve-2010-0213).

*13 DHCP: Fencepost error on zero-length client identifier (http://www.isc.org/software/dhcp/advisories/cve-2010-2156).

*14 Cisco Security Advisory: Cisco IOS XR Software Border Gateway Protocol Vulnerability (http://www.cisco.com/en/US/products/products_security_advisory09186a0080b4411f.shtml).

*15 Cisco Security Advisory: Summary of Cisco IOS Software Bundled Advisories, September 22, 2010 (http://www.cisco.com/warp/public/707/cisco-sa-20100922-bundle.shtml).

*16 About the security content of iOS 4.1 for iPhone and iPod touch (http://support.apple.com/kb/HT4334).

*17 There are many detailed reports related to this malware, for example the following commentary from Nippon CSIRT Association. About the Stuxnet malware (http://www.nca.gr.jp/2010/stuxnet/index.html) (in Japanese).

*18 A form of SCADA (Supervisory Control And Data Acquisition). This is a solution for monitoring systems and controlling processes via computer that is generally used at facilities such as factories.

*19 Fraudulent incoming SIP 24 (http://jvnrss.ise.chuo-u.ac.jp/csn/index.cgi?p=%C9%D4%C0%B5%A4%CASIP%C3%E5%BF%AE+24) (in Japanese). cNotes provides SIP observation data on an irregular basis.

*20 Details of this vulnerability can be found in the following official blog post. Twitter blog - All about the "onMouseOver" incident (http://blog.twitter.com/2010/09/all-about-onmouseover-incident.html).

*21 Details of this incident can be found in the following F-Secure blog post. Worms Loose on Twitter.com (http://www.f-secure.com/weblog/archives/00002034.html).

*22 This incident is also detailed in the following Trend Micro blog post. Adobe zero-day attacks and Web-based threats via ad distribution systems - looking back on threat trends for September 2010 (http://blog.trendmicro.co.jp/archives/3700) (in Japanese).

*23 Information about DNSSEC for the Root Zone (http://www.root-dnssec.org/2010/07/16/status-update-2010-07-16/).

*24 JPRS Plans to Implement DNSSEC in JP Domain Name Services in January 2011 (http://jprs.co.jp/en/topics/2010/100728.html).

*25 Fixes regarding the TLS renegotiation feature are included in the following program update. Microsoft Security Bulletin MS10-049 - Critical: Vulnerabilities in SChannel could allow Remote Code Execution (980436) (http://www.microsoft.com/technet/security/bulletin/ms10-049.mspx). This issue is explained in Vol.6 of this report under "1.4.2 MITM Attacks Using a Vulnerability in the SSL and TLS Renegotiation" (http://www.iij.ad.jp/en/development/iir/pdf/iir_vol06_EN.pdf).

*26 MOAUB (Month of Abysssec Undisclosed Bugs). Details of the vulnerabilities reported through this initiative can be found on the Abysssec Security Research blog. MOAUB - Day by Day (http://www.abysssec.com/blog/2010/09/moaub-1/).

1.3 Incident Survey

Of incidents occurring on the Internet, IIJ focuses on those types of incidents that have infrastructure-wide effects, continually conducting research and engaging in countermeasures. In this section, we provide a summary of our survey and analysis results related to the circumstances of DDoS attacks, malware infections over networks, and SQL injections on Web servers.

1.3.1 DDoS Attacks

Today, DDoS attacks on corporate servers are almost a daily occurrence. The methods involved in DDoS attacks vary widely. Generally, however, these attacks are not the type that utilize 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 IIJ DDoS Defense Service between July 1 and September 30, 2010.

This information shows traffic anomalies judged to be attacks based on IIJ DDoS Defense Service standards.

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*27, attacks on servers*28, and compound attacks (several types of attacks on a single target conducted at the same time).

During the three months under study, IIJ dealt with 622 DDoS attacks. This averages to 6.76 attacks per day, which is three times the average daily number of attacks observed during the period for our prior report. Bandwidth capacity attacks accounted for 1% of all incidents, server attacks accounted for 72% of all incidents, and compound attacks accounted for the remaining 27%. This is due to a drastic increase in attacks on multiple Web servers that occurred over the period of September 10 to September 30, which made up 46% of the total.

The largest attack observed during the period under study was classified as a bandwidth capacity attack, and resulted in 1.4Gbps of bandwidth using up to 275,000pps packets. Of all attacks, 66% ended within 30 minutes of commencement, while 20% lasted between 30 minutes and 24 hours. The longest attack continued for 12 days (291 hours), and consisted of a compound attack that resulted in 670Mbps of bandwidth using up to 120,000pps packets.

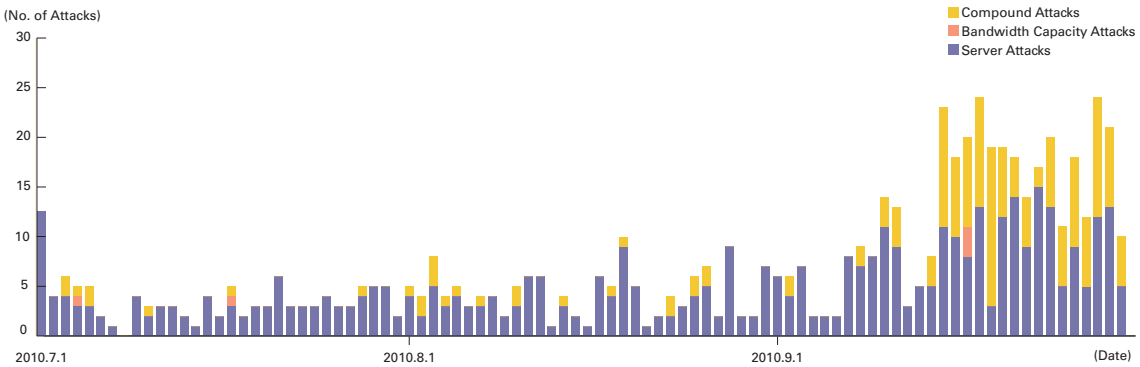

Figure 2: Trends in DDoS Attacks

*27 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.

*28 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.

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*29 and botnet*30 usage as the method for conducting DDoS attacks.

Backscatter Observations

Next we present our observations of DDoS backscatter using the honeypots*31 set up by the MITF, a malware activity observation project operated by IIJ*32. By monitoring backscatter it is possible to detect certain types of DDoS attacks occurring on external networks as a third party without any interposition.

Figure 3 shows trends in packet numbers by port for the backscatter observed between July 1 and September 30, 2010, and Figure 4 shows the sender's IP addresses classified by country.

The port most commonly targeted by the DDoS attacks observed was the 80/TCP port used for Web services, accounting for 58.4% of the total during this period. Attacks on 3389/TCP used for remote desktop were also observed. Additionally, many attacks were observed on ports not used by common applications, such as 5218/TCP and 5224/TCP. Looking at the origin of backscatter thought to indicate IP addresses targeted by DDoS attacks by country in Figure 4, China and the United States accounted for large proportions at 44.8% and 29.5%, respectively, and Japan made up 2.1% of the total. The particularly large number of backscatter packets observed targeting 5224/TCP on August 20 and 5218/TCP on September 19 all show a single IP address in China as the attack target.

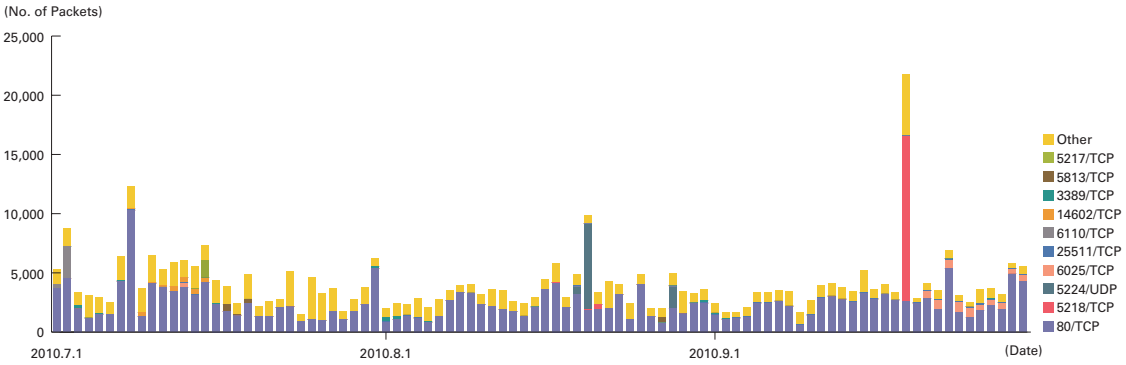

Figure 3: DDoS Attack Backscatter Observations (Observed Packets, Trends by Port)

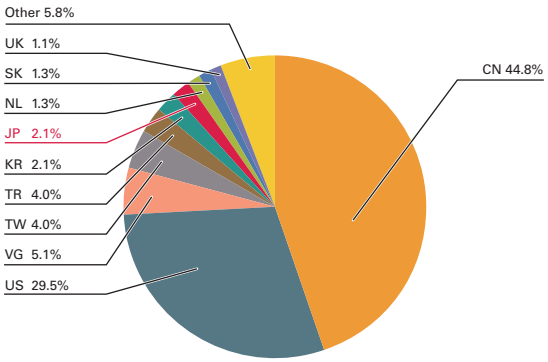

Figure 4: DDoS Attack Target Distribution According to Backscatter Observations (by Country, Entire Period under Study)

*29 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 in order to pretend that the attack is coming from a different location, or from a large number of individuals.

*30 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."

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

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

1.3.2 Malware Activities

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

■ Status of Random Communications

Figure 5 shows trends in the total volumes of communications coming into the honeypots (incoming packets) between July 1 and September 30, 2010. Figure 6 shows the distribution of sender's IP addresses by country. The MITF has set up numerous honeypots for the purpose of observation. We have taken the average per honeypot, showing the trends for incoming packet types (top ten) over the entire period subject to study.

Much of the communications arriving at the honeypots demonstrated scanning behavior targeting TCP ports utilized by Microsoft operating systems. We also observed scanning behavior for 1433/TCP used by Microsoft's SQL Server and 23/TCP used for telnet. Additionally, attacks were observed on ports not used by common applications, such as 5121/TCP, 31795/TCP, 23502/TCP, and 9415/TCP. Looking at the overall sender distribution by country in Figure 6, we see that attacks sourced to Japan at 30.4%, China at 15.9%, and Taiwan at 6.0% were comparatively higher than the rest.

■ Malware Network Activity

Next, we will take a look into the malware activity observed by the MITF. Figure 7 shows trends in the total number of malware specimens acquired during the period under study. Figure 8 shows the distribution of the specimen

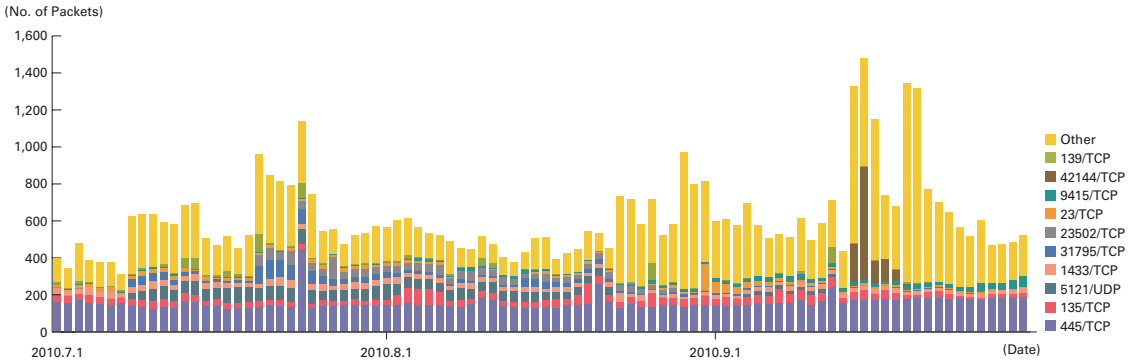

Figure 5: Communications Arriving at Honeypots (by Date, by Target Port, per Honeypot)

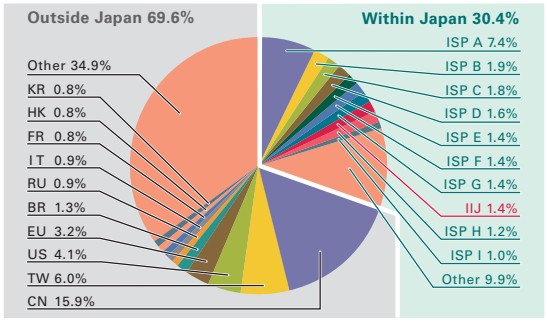

Figure 6: Sender Distribution (by Country, Entire Period under Study)

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

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

acquisition source for malware. In Figure 7, the trends in the number of acquired specimens show the total number of specimens acquired per day*35, while the number of unique specimens is the number of specimen variants categorized according to their digest of a hash function*36.

On average, 371 specimens were acquired per day during the period under study, representing 41 different malware variants. Statistics in our prior report show that the average daily total for acquired specimens was 378, with 32 different variants. During the current period the total number of specimens acquired decreased slightly, but the number of different variants rose compared to the previous period. The sharp drop in total specimens acquired after September 19 is due to the activity of Sdbot and its variants ceasing worldwide. The reason for this suspension of Sdbot activity is not known.

The distribution of specimens according to source country in Figure 8 has Japan at 38.7%, with other countries accounting for the 61.3% balance. Taiwan was at 47.8%, maintaining the large percentage that it held during the previous period. This is due to the high level of activity shown by Sdbot and its variants in Taiwan, but as with other countries this activity ceased after September 19.

The MITF prepares analytical environments for malware, conducting its own analyses of acquired specimens. During the current period under observation 14.0% of the malware specimens acquired were worms, 84.8% were bots, and 1.2% were downloaders. In addition, the MITF confirmed the presence of 26 botnet C&C servers*37 and 276 malware distribution sites. The number of malware distribution sites detected increased dramatically due to the increase in specimens accessing multiple distribution sites, which had dropped in the previous period under study.

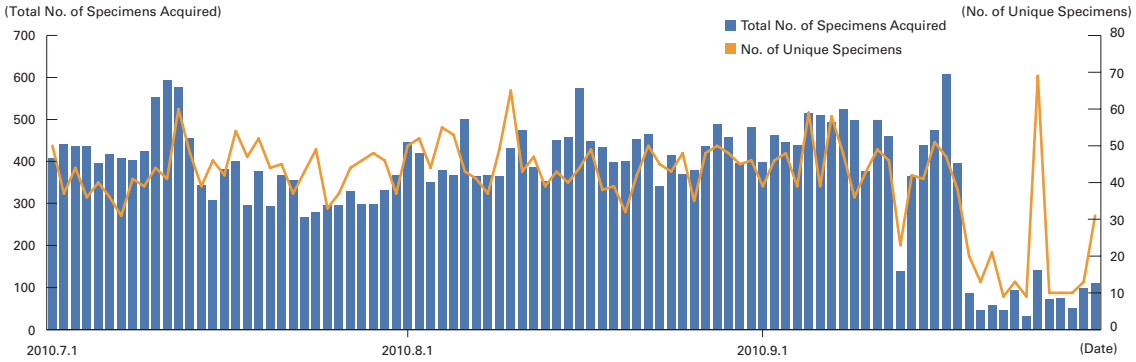

Figure 7: Trends in the Number of Malware Specimens Acquired (Total Number, Number of Unique Specimens)

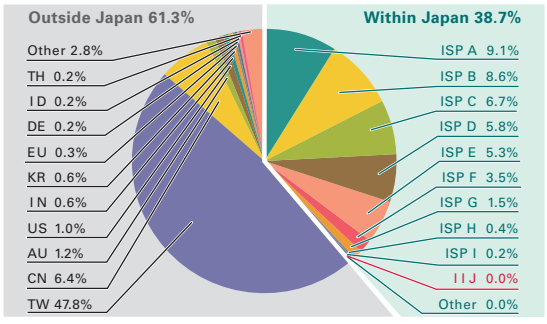

Figure 8: Distribution of Acquired Specimens by Source (by Country, Entire Period under Study)

*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." A server that provides commands to a botnet consisting of a large number of bots.

1.3.3 SQL Injection Attacks

Of the types of different Web server attacks, IIJ conducts ongoing surveys related to SQL injection attacks*38. SQL injection attacks have flared up in frequency numerous times in the past, remaining one of the major topics in the 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 9 shows trends of the numbers of SQL injection attacks against Web servers detected between July 1 and September 30, 2010. Figure 10 shows the distribution of attacks according to source. These are a summary of attacks detected by signatures on the IIJ Managed IPS Service.

Japan was the source for 40.0% of attacks observed, while China and the United States accounted for 36.7% and 7.1%, respectively, with other countries following in order. There was very little change from the previous period in the status of SQL injection attacks against Web servers. However, as attacks mainly from China against specific targets attempting to gain access privileges on SQL servers occurred on September 30, the percentage of attacks accounted for by China has increased.

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

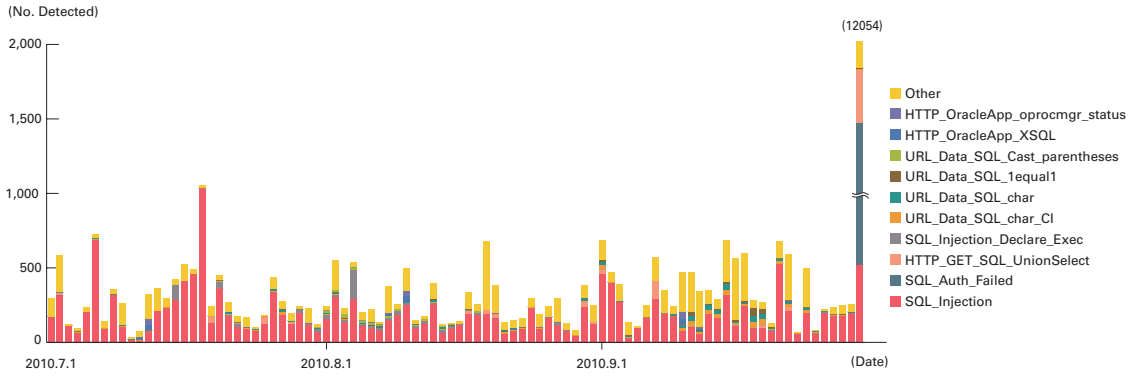

Figure 9: Trends in SQL Injection Attacks (by Day, by Attack Type)

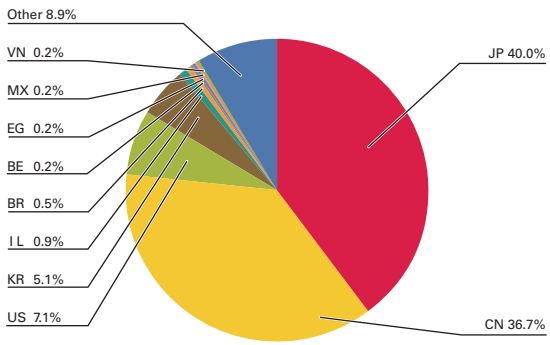

Figure 10: Distribution of SQL Injection Attacks by Source (by Country, Entire Period under Study)

*38 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.

1.4 Focused Research

Incidents occurring over the Internet change in type and scope almost from one minute to the next. Accordingly, IIJ works toward taking countermeasures by continuing to perform independent surveys and analyses. Here we will present information from the surveys we have undertaken during this period regarding preparations to be made for DDoS attacks on small-scale systems and security for shared systems, and also provide an overview of digital forensics.

1.4.1 Preparing for DDoS Attacks on Small-Scale Systems

As shown in "1.3.1 DDoS Attacks," multiple DDoS attacks on Web servers took place during September for the period covered in this report. These attacks targeted the Web servers of both Japanese public agencies and private-sector businesses. DDoS attacks are generally carried out as a form of protest or a personal statement to the owner of the server that is targeted. In recent years there have also been cases in which DDoS attacks are used in blackmail attempts to extort money. Technology such as dedicated attack tools and botnets are used in current DDoS attacks, and some people will even carry out attacks on behalf of others. In other words, those intending to carry out an attack can do so comparatively easily without specialist knowledge or technology. This means there is a chance that any server exposed to the Internet could be the target of a DDoS attack regardless of its scale. Here we examine preparations to be made for DDoS attacks on small-scale systems.

DDoS attacks include those that overload the server itself directly, and those that flood the lines being used by the server with communications. Both of these types of attack cause servers connected to the Internet to be suddenly rendered inoperative. When evaluating countermeasures that will be effective when such an attack occurs, it is necessary to make preparations such as establishing a countermeasure policy based on the server's level of importance, improving the server's tolerance, building a system for abnormal behavior detection, and requesting the cooperation of other organizations.

■ Establishing a Countermeasure Policy

Servers targeted by a DDoS attack will no longer be able to carry out their intended role, which is a threat to availability. Consequently, you should first clarify how business would be affected if a server under examination was to suspend its operations. Once this is done, a target should be set for recovery of server functionality. For example, it is necessary to consider factors such as whether a complete suspension would be problematic for the business, whether it is possible to limit the scope of communications to a certain domain (such as within Japan or business partners only), and whether the quality of communications can be lowered (limiting connections from a certain address or applying bandwidth control across the board). Look into introducing a dedicated DDoS countermeasure device or using a countermeasure service for servers that require high availability.

■ Improving Server Tolerance

Communication lines may become flooded and servers overloaded when you are targeted by a DDoS attack. This means it may be difficult to confirm the communications or operating status of unprotected servers. When installing servers, it is necessary to evaluate the processing ability required for normal business operations and incorporate

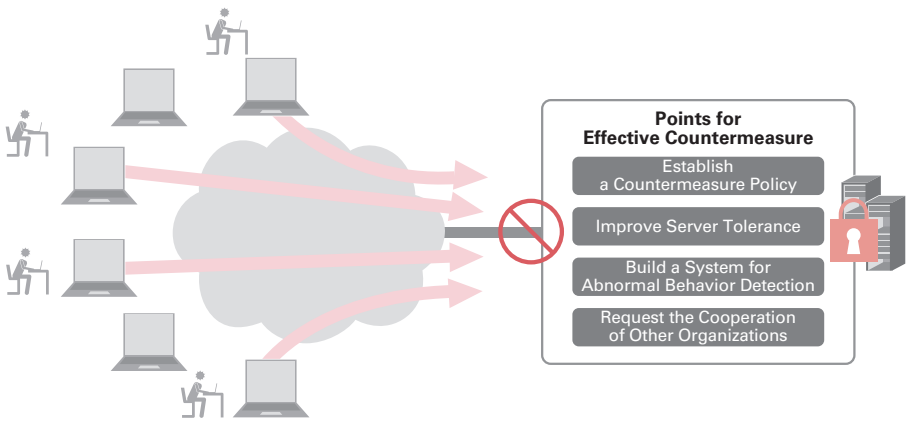

Figure 11: Preparing for DDoS Attacks

a certain surplus over and above this. It is also important to look into implementing DDoS countermeasures and resource management options for the server's OS and applications. For example, Linux incorporates the SYN cookies function*39 that protects against SYN flood attacks, and a function that limits connections for each application*40. Apache HTTP servers also offer functions for limiting elements such as simultaneous connections in their settings*41 or by adding external modules*42. By combining hardware performance with OS and application functions such as these, it is possible to increase a server's tolerance of DDoS attacks.

■ Building a System for Abnormal Behavior Detection

When trying to ascertain the situation after being targeted by a DDoS attack, it can take an extremely long time to complete analysis due to suitable logs not being acquired on a daily basis or the large volume of log data that must be processed. Consequently, it is crucial to record logs appropriately and confirm communications status on a regular basis to prepare a system that detects DDoS attacks as an abnormality. In addition to server logs, referring to data regarding communications from SNMP or Netflow will make it easier to grasp the situation. It is also necessary to prepare for handling logs containing large volumes of data. For example, by setting up a log server separate to the Web servers it is possible to analyze records even when the Web servers are overloaded. Additionally, by preparing scripts for extracting a summary of large logs in advance, it is possible to understand the situation swiftly when abnormal behavior is detected.

■ Requesting the Cooperation of Other Organizations

Sometimes it may not be possible for a targeted server to deal with attacks that flood lines with communications or spoofed IP addresses. In such cases it may be necessary to request help from external organizations such as ISPs, security vendors, or CSIRTs. When doing this you will need to disclose what you have already learned about the situation. Additionally, in many cases the organization from which help is requested will not be able to deal with the threat alone, and it will be necessary to give permission for them to share information about the attack with other organizations (such as the attacker's ISP). By determining the attack information that can be disclosed in advance, such as IP address, attack details, and communication patterns, it will be possible for other organizations to implement countermeasures swiftly. Documents such as the JPCERT/CC incident report form*43 serve as useful reference points for this kind of information.

■ Summary

In this section we have covered a number of points that should be prepared in advance for small-scale systems that may be targeted by DDoS attacks. As shown in the backscatter observations in "1.3.1 DDoS Attacks," DDoS attacks on servers not providing Web content have been observed, and there is a need to prepare for sudden DDoS attacks on all kinds of servers. DDoS attacks embody a strong message from the attacker, and their occurrence can be anticipated to a certain extent. For this reason, paying attention to world trends and news related to your company such as information regarding organizations you belong to and identifying the precursors to conflict at an early stage can serve as helpful preparations against DDoS attacks*44.

1.4.2 Shared System Security

Recently, the full-scale use of cloud computing has begun to accelerate. The cloud makes it possible to use a variety of system resources at low cost by sharing them, but the unique security issues faced by shared systems are of concern. Here we examine the threats that occur in the cloud due to shared system resources as well as their countermeasures.

*39 See Daniel J. Bernstein's explanation of SYN cookies (http://cr.yp.to/syncookies.html) for more details. IETF's RFC4987 "TCP SYN Flooding Attacks and Common Mitigations" (http://www.ietf.org/rfc/rfc4987.txt) also provides a summary covering the principles behind SYN flood attacks as well as countermeasure technologies.

*40 iptables provides modules such as limit and connlimit. For example, using the iptables limit module to limit syn packets makes it possible to set a cap on the number of new connections that an application can process.

*41 It is possible to use the MaxClients setting to limit the number of simultaneous connections. You can also contain the resources consumed by an attack by adjusting settings such as Timeout, KeepAlive, KeepAliveTimeout, and MaxKeepAliveRequests.

*42 A large number of external modules exist for Apache. One example, mod_limitipconn (http://dominia.org/djao/limitipconn2.html) makes it possible to limit the number of simultaneous connections from a single IP address.

*43 See the following JPCERT Coordination Center page on incident reporting for more details (http://www.jpcert.or.jp/english/ir/form.html).

*44 Other useful information on DDoS attacks includes VeriSign Inc.'s "DDoS Mitigation – Best Practices for a Rapidly Changing Threat Landscape Whitepaper" (user registration required) (http://www.verisign.com/forms/ddosbestpracticeswp.html?toc=MYUM9-0000-02-00).

Issues with Shared Resources

System resources are shared in a multi-tenant cloud, with the partitioning of resources between users handled logically by software, etc. For this reason it is crucial this logical resource partitioning is handled appropriately, as breaking the boundaries of a partition would represent a security threat for users. Let us examine the potential threats that could actually occur. The CSA (Cloud Security Alliance)*45 lists seven items as threats to cloud computing in "Top Threats to Cloud Computing V1.0"*46. These include "Shared Technology Issues," which covers incidents and impact from the improper logical partitioning of shared resources such as CPU and GPU. Shared resources have also been brought up as a threat unique to cloud computing in many articles other than the CSA report. As resources such as communication lines, communication devices, and storage are also shared in the cloud, these must also be considered. Here we look at concrete examples of threats while considering cloud system architecture.

Architecture of Cloud Infrastructure

The architecture of a cloud is generally not made public. For this reason, we assume a cloud composed of generic equipment, and use a cloud system comprised of the devices shown in Figure 12 as an example to evaluate threats to cloud computing.

This cloud uses a router or similar device to connect to the Internet. Users access VM (virtual machines) on the cloud via the Internet (the red arrow in Figure 12). The physical machines running these virtual machines also accommodate the virtual machines of other users, so the physical resources are shared. The physical machines have multiple Ethernet ports for providing service, including those for connecting to the Internet and those providing storage services via an Ethernet connection to a storage network (IP-SAN).

Users access VMs using the route shown by the red line in the figure. Users are generally not aware that they are working on a complex device layout when using the cloud. A shared environment that is not visible to its users may have inherent security issues.

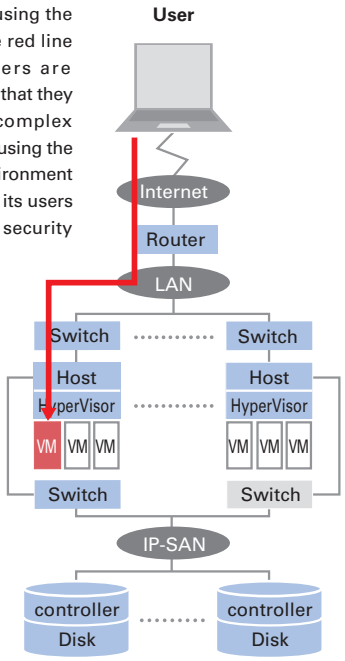

Figure 12: Sample Cloud System Architecture

Cloud Environment Threats and their Countermeasures

Sharing network resources is not unique to cloud environments, as the Internet itself and rental server environments also involve multiple users sharing a network. However, in this sample cloud there is the possibility that virtual machines with different security policies will occupy that same L2 segment. A virtual machine could be hacked and hijacked, presenting the risk of communications being blocked or intercepted via methods such as ARP poisoning*47. To counter threats such as these it is necessary to implement measures for preventing the spoofing of the VLAN ID or MAC address in hypervisor or the connected switch.

In this sample cloud, storage is also shared through virtualization. When technology such as IP-SAN (iSCSI) is used, storage is connected via Ethernet, making it possible to attack a storage network by artificially generating a fake Ethernet frame. When a storage controller exists on a network location reachable by virtual machines, attacks against this controller are also possible. Additionally, when the ID (IQN*48) for storage virtualization is falsified, there is a risk that data areas that should be partitioned and invisible could be

*45 CSA (Cloud Security Alliance) is an organization established in 2008 to promote best practices for cloud security (http://www.cloudsecurityalliance.org/). The Cloud Security Alliance Japan Chapter came into being in June 2010 as the Japan branch of the organization (http://www.cloudsecurityalliance.jp/) (in Japanese).

*46 In "Top Threats to Cloud Computing V1.0," the CSA documents threats that are typical to cloud computing as well as their countermeasures (http://www.cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf).

*47 ARP poisoning is an attack that involves sending spoofed ARP packets to a network to intercept communications from other hosts.

*48 An abbreviation of "iSCSI Qualified Name." This is the name used to identify iSCSI nodes on a network.

viewed and connected to. To counter threats such as these it is necessary to implement access controls or spoofing countermeasures in hypervisor or the switch in the same way as with network threats.

Additionally, when basic functions such as firewalls are provided as part of a service, threats vary depending on the form in which they are provided. Figure 13 shows sample firewall placements. Pattern 1 represents the method in which a firewall is provided as a hypervisor function. In this case firewall security is guaranteed by the service provider. In pattern 2 a software firewall solution is applied by installing an OS on the hypervisor in the same way as the virtual machines provided to users. This would mean that when vulnerabilities exist in the hypervisor, the firewall would also be affected. In pattern 3 a dedicated device is prepared and network traffic relayed through it using VLAN and routing. This method replicates existing models, but has a higher cost. In pattern 4 the firewall is implemented on the user's OS, but this presents the risk of settings being changed by hackers. Pattern 5 represents a method often used for the Web and email, in which an application's proxy function is used as SaaS. When services are offered by several different providers, regulation between services is the responsibility of the user. These examples demonstrate the need to be aware that different points must be considered depending on the form of service provided.

■ Summary

The items explained here are not new concepts, and under normal conditions as long as the service provider is aware of each threat and takes appropriate countermeasures, there should be no risk of issues occurring.

It is helpful for users to be aware of internal architecture to confirm security, but in many cases this kind of information is not disclosed when using a service. When using a cloud, it is important for users to assess their scope of responsibility with regard to the services they use and implement the necessary countermeasures, while also reaching agreement regarding security measures with the service provider.

1.4.3 An Overview of Digital Forensics

With the spread of IT much of the information retained by companies and individuals is saved and accumulated as digital data. Because of this, cases in which digital data is used in incident responses or as evidence in trials are on the rise. Digital data is more easily changed or destroyed than physical media, so those who investigate such matters must have access to appropriate technology. Here we will explain the digital forensics techniques that are used to examine digital data.

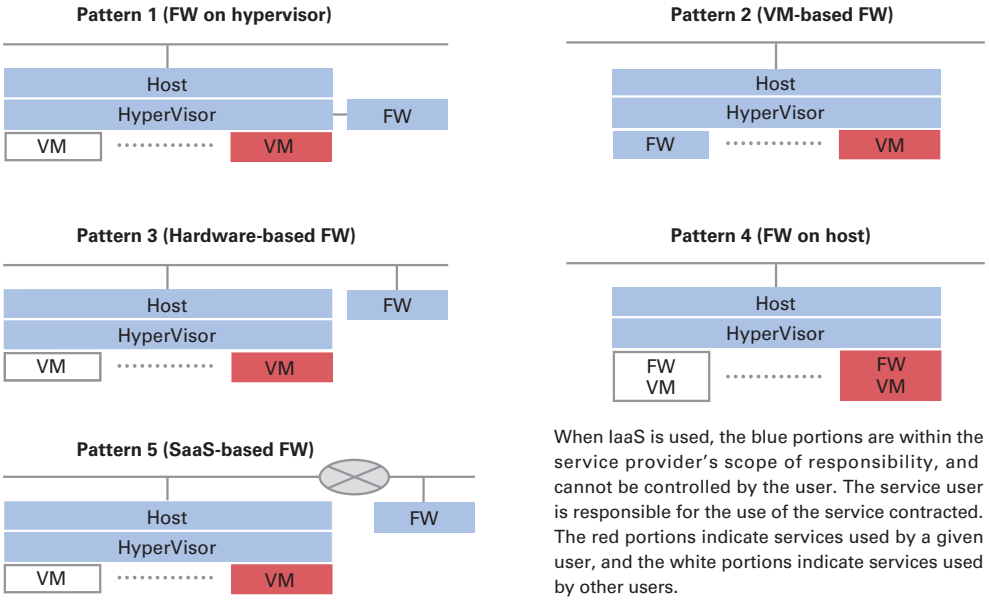

Figure 13: Sample Firewall Placements

Digital forensics is a technology used mostly in corporate environments*49 in situations such as incident responses investigating unauthorized access or when presenting digital data for a trial. To categorize digital forensics from the perspective of the subject matter being analyzed, it includes computer forensics for analyzing computers, network forensics*50 for analyzing packets sent over a network, and mobile device forensics*51 for analyzing mobile devices such as mobile phones.

Digital forensics is carried out by acquiring, analyzing, and reporting data in that order. Apart from certain exceptions, digital forensics is generally not carried out by using the original digital data for analysis. Engineers carrying out digital forensics first acquire the data in question. Next, the acquired data is analyzed using appropriate tools, and the series of events related to the incident are reconstructed. Last of all, the established facts are put together in a report based on the results of the analysis.

■ Computer Forensics

Next, we will cover the acquisition and analysis process for computer forensics, which is frequently used.

Computers are comprised of components such as the CPU, memory, and hard disks. Data on hard disks or CD-ROM media is non-volatile and remains intact even when the computer is powered down. Data in the CPU and memory is volatile, and is lost when the power is turned off. In RFC3227 "Guidelines for Evidence Collection and Archiving"*52 it is recommended that when acquiring data you should proceed from the volatile to the less volatile. Accordingly, when the computer being examined is a server in an online state (powered on and running), it is best to first collect volatile data such as that found in memory before collecting data on the hard disk or other backup media.

Methods for collecting volatile data include executing a volatile data gathering toolkit on the machine in question, and acquiring a memory image for later analysis. An example of a volatile data collection toolkit is Sysinternals Suite*53.

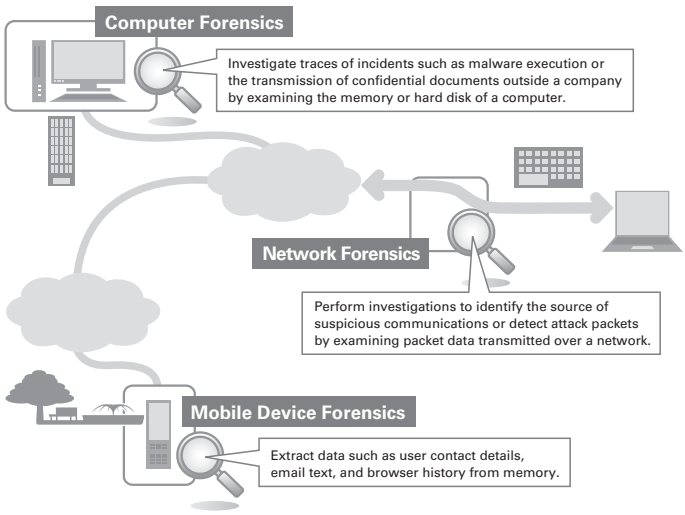

Figure 14: Forensics Overview

*49 NPO The Institute of Digital Forensics defines digital forensics as "a series of scientific investigation methods and technologies for acquiring evidence, investigating, and analyzing electromagnetic records for incident response and legal disputes/litigation, and the analysis and data collection related to the alteration or damage of electromagnetic records." (http://www.digitalforensic.jp/wdfitm/wdf.html) (in Japanese).

*50 Techniques for investigating and analyzing computers remotely are also sometimes referred to as network forensics.

*51 Mobile device forensics was seldom implemented outside law enforcement agencies in Japan because the forensic tools used in Europe and America do not support Japanese mobile phone specifications. However, tools compatible with the smartphones that are growing in popularity recently have started to become available for use in Japan. For example, Oxygen Software's Oxygen Forensic Suite tool for mobile device analysis is now used by the Cyber Defense Institute Inc. (http://www.cyberdefense.jp/company_profile/prerelse10001.html) (in Japanese).

*52 RFC3227 "Guidelines for Evidence Collection and Archiving" (http://www.ietf.org/rfc/rfc3227.txt) contains evidence collection procedures and cautionary notes for when a security incident occurs.

*53 Sysinternals Suite includes a wide variety of programs, such as PsList for showing information about the processes being executed and TcpView for showing TCP connection status (http://technet.microsoft.com/en-us/sysinternals/bb842062.aspx).

This group of tools is executed on the target machine, with results output as a text file or similar format. Meanwhile, a memory image refers to a raw dump of the contents of memory in binary file format. To acquire a memory image a tool is executed on the target machine to dump the memory, and this memory image is later transferred to another machine for analysis to extract the volatile information. Memory images must always be acquired before executing a volatile data collection toolkit*54.

Methods for acquiring non-volatile data include online and offline acquisition (with the computer in a powered down state). Data may be saved as physical disk images for acquiring each physical disk, or as logical disk images for acquiring each logical volume. In general, non-volatile data is acquired offline. However, when it is more convenient to acquire data as logical volumes, such as when disks are encrypted or in a RAID configuration, acquisition may also be carried out online. A hash value is calculated for the acquired data at the time of acquisition to ensure the integrity of the data after it is acquired. This makes it possible to detect whether data has been changed or falsified after acquisition.

By analyzing acquired volatile data it is possible to obtain information such as the user that was logged in, the communications status and start times of the running processes, the files and shared libraries accessed by processes, the ARP table, the routing table, and the DNS cache. It may also be possible to acquire information such as processes previously executed and communications that had already been closed by analyzing memory image files.

A broad array of results can be acquired through analysis of non-volatile data. Because related executable files and logs are often deleted during a hacking or malware infection incident to conceal the attack, deleted files are recovered and log fragments from unallocated disk space examined. Furthermore, if it is possible to detect operations or files related to an incident from system and event logs or the registry, you can estimate the time an incident occurred and the scope of damages by examining other events close to the time the operations were carried out or the time stamp (data indicating the time a file was created or updated) of corresponding files.

Also, by searching for keywords in acquired data, whether volatile or non-volatile, it is possible to confirm traces of data communications to external parties, and check whether other systems have been accessed*55.

Issues with Digital Forensics

The following issues are present in digital forensics.

- Consolidating logs from multiple sources
- Dealing with anti-forensics
- Dealing with large volumes of media and encryption

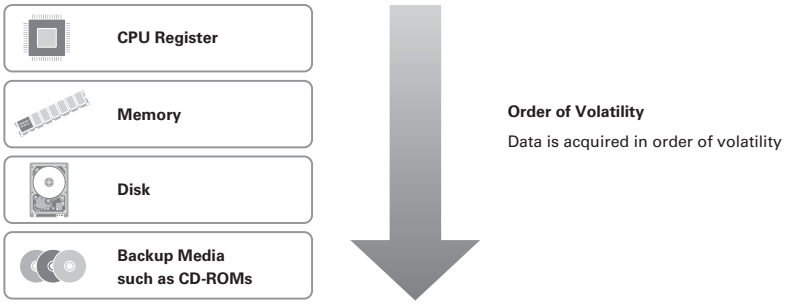

Figure 15: Order of Volatility

*54 A volatile data collection toolkit is composed of a large number of programs that may influence the content of a memory image due to actions such as a swap occurring when they are executed.

*55 Examples of the tools for conducting data analysis indicated here include EnCase (http://www.guidancesoftware.com/), FTK (http://www.accessdata.com/), and TSK (http://www.sleuthkit.org/).

An example of a case affected by the issues with consolidating logs from multiple sources is the examination of the extent that confidential information has spread in an information leak. When conducting such an examination, network forensics must be carried out in addition to computer forensics. Specifically, it is necessary to compare the logs and packet data from network devices such as firewalls, routers, and IDS with the data from the computer. However, when the logs for each device have different formats or time settings this comparison work becomes an enormous task.

Anti-forensics are techniques for evading detection by digital forensics. For example, to evade techniques that examine related files based on file system time stamps, some malware sets a random time in the time data of files it creates. To deal with this kind of anti-forensics, it is necessary to utilize time stamps other than those set on the file system*56.

The size of media to be acquired is also growing year-on-year. To acquire data in the shortest time possible the only option is to make improvements to the performance of acquiring devices and software. However, when there is not enough time, one technique that can be utilized is previewing data before it is acquired (examining the given media directly by mounting it with read-only access). It would also be desirable to automate processes to analyze large volumes of media in an efficient manner. There have recently been many cases in which the target media has been encrypted. In this case it is necessary to either acquire data online, or decrypt the data once it has been acquired offline.

Summary

Using digital forensics in incident responses makes it possible to examine the cause of an incident as well as the extent of its impact without overlooking any key data. IIJ will continue to look into ways of resolving the issues faced by digital forensics while constantly evaluating the latest technology trends and applying these findings to our response methods.

1.5 Conclusion

This report has provided a summary of security incidents to which IIJ has responded. In this report we examined the preparations necessary for incident response by looking at preparations to be made for DDoS attacks on small-scale systems and security for shared systems such as a cloud computing, as well as giving an overview of digital forensics.

By identifying and publicizing incidents and associated responses in reports such as this, IIJ will continue to inform the public about the dangers of Internet usage, providing the necessary countermeasures to allow the safe and secure use of the Internet.

Authors:
Mamoru Saito
Manager of the Office of Emergency Response and Clearinghouse for Security Information, IIJ Service Division. After working in security services development for enterprise customers, Mr. Saito became the representative of the IIJ Group emergency response team, IIJ-SECT in 2001, participating in FIRST, an international group of CSIRTs. Mr. Saito serves as a steering committee member of several industry groups, including Telecom-ISAC Japan, Nippon CSIRT Association, Information Security Operation provider Group Japan, the Web Malware Mitigate Community, and others. He is also active in organizations such as the Engineers SWG of the Anti-Child Pornography WG in the Association for Promoting the Creation of a Safe Internet, and the IPA Conference for Denial of Service Attack Countermeasures.

Hirohide Tsuchiya (1.2 Incident Summary)
Hirohide Tsuchiya, Hiroshi Suzuki, Tadaaki Nagao (1.3 Incident Survey)
Mamoru Saito, Hirohide Tsuchiya (1.4.1 Preparing for DDoS Attacks on Small-Scale Systems)
Masahiko Kato (1.4.2 Shared System Security)
Takahiro Haruyama (1.4.3 An Overview of Digital Forensics)
Office of Emergency Response and Clearinghouse for Security Information, IIJ Service Division

Contributors:
Yuji Suga, Hiroaki Yoshikawa, Tadashi Kobayashi, Seigo Saito
Office of Emergency Response and Clearinghouse for Security Information, IIJ Service Division

*56 Examples of other time stamps include data recorded within files, the data for an original file found in a shortcut file, and registry key information. For malware infections made through Web browsers or emails, browser history and the times email has been sent or received can also be of help.