Understanding Issues in Cloud Forensics: Two Hypothetical Case Studies

Josiah Dykstra, Alan T. Sherman
Department of Computer Science and Electrical Engineering
University of Maryland, Baltimore County
{dykstra, sherman}@umbc.edu

Abstract

The inevitable vulnerabilities and criminal targeting of cloud environments demand an understanding of how digital forensic investigations of the cloud can be accomplished. We present two hypothetical case studies of cloud crimes; child pornography being hosted in the cloud, and a compromised cloud-based website. Our cases highlight shortcomings of current forensic practices and laws. We describe significant challenges with cloud forensics, including forensic acquisition, evidence preservation and chain of custody, and open problems for continued research.

1 Introduction

Crime committed using cloud computing resources and against cloud infrastructures is inevitable. Though real incidents have already taken place against cloud providers including Google, an absence of documentation indicates that no crimes using the cloud or targeting it directly have been publicized nor litigated thus far. Forensic investigators must understand that current tools and techniques are inadequate in the cloud environment where acquisition, examination and analysis will be in practice executed very differently than is done today. To illustrate these issues, we fabricate two hypothetical crimes and deconstruct the forensic investigation against them.

Companies are embracing cloud technology to offload some of the cost, upkeep, and growth of equipment that they would otherwise have purchased themselves. Cloud infrastructure, with exceptional bandwidth, storage and computing power, offers an attractive prize for hackers. While many people have lamented how the users of the cloud and their data are protected, few of these discussions have considered the difficulty of responding to security breaches, including forensics and criminal prosecution.

In this article, we consider the investigative response and forensic process of two hypothetical, but plausible, case studies of crimes tied to cloud computing. In Section 2, we present previous and related work. In Section 3 we discuss the applicability of forensic frameworks. Section 4 contains our case studies. The first explores a case of child pornography in the cloud, and the trouble with both acquiring and analyzing data. The second case study deals with the cloud as the target of a crime, and the complex issues of chain of custody and trust. We examine issues of attribution, forensic integrity and chain of custody in Section 5, and we conclude in Section 6.

2 Previous Work

Despite significant research in digital forensics, little has been written about the applicability of forensics to cloud computing environments. Furthermore, no case law exists on which to extrapolate the desire of the courts on the matter. Garfinkel recently suggested that "cloud computing in particular may make it impossible to perform basic forensic steps of data preservation and isolation on systems of forensic interest" (Garfinkel 2010). In one of the only published books on cloud forensics, the subject is approached as a matter of network forensics combined with remote disk forensics (Lillard 2010). While legal complications are introduced, including cloud-based evidence admissibility, no solutions are presented. Wolthusen identified some research challenges, including "discovery of computation structure," "attribution of data," "stability of evidence," and "presentation and visualization of evidence" (Wolthusen 2009). In 2009, researchers at UC San Diego demonstrated that it was possible to locate a particular virtual machine (VM) in Amazon Elastic Compute Cloud (EC2) and mount side-channel attacks by co-locating a new VM with the target (Ristenpart 2009).

In 2009, Google and 34 other companies were hacked and infected with data-stealing malware. While the attack at Google involved Gmail, a cloud-based email service, the vulnerabilities and exploits were end-user based and not an attack on the cloud (Symantec 2010). Using Amazon EC2, researchers recently demonstrated how to crack passwords quickly and cheaply, a potentially criminal activity (Bagh 2011). In 2010, presenters at the DEFCON Conference used EC2 to launch a demonstration denial of service against a small network (Lemos 2010). In the investigation of individual users, cloud providers have begun to offer services that aid law enforcement. For example, Facebook gives a user the option to download their entire personal profile and history (Facebook 2011). However promising this may be for an investigator, these data cannot be said to be forensi-

cally sound. Guidance Software, the maker of EnCase, has produced a training video showing how to recover and analyze Facebook chat artifacts from a local hard drive (Guidance 2009).

Lawyers and computer scientists alike have expressed views about remote forensics, a field that shares an important similarity to the cloud. Schwerha and Inch (Schwerha and Inch 2008) list remote forensic software and survey legal analysis and case law. They undertook no application to cloud computing. Law professor Orin Kerr has written extensively on the applicability of the Fourth Amendment to electronic evidence and the Internet (Kerr 2009). His suggestions on search warrant language for shared resources are apropos to cloud forensic research. In Australia, lawmakers are already being made aware that current law enforcement is not equipped to investigate attacks on cloud services (Choo 2009).

3 Frameworks

To frame the approach of forensic investigation of any environment, including the cloud, it is helpful to have a procedure that guides the activity. The cloud environment does not affect the need for a framework, and does not inherently demand a new one. Frameworks for the digital forensic investigation are plentiful: at least 14 have been published since 1995 (Selamat et al. 2008). Digital forensic labs often choose a combination of approaches, or develop their own process that considers their particular personnel, workload, and budget. The generality of many investigative frameworks makes them applicable under many circumstances and irrespective of technology. While there is hardly a generic computer forensic case that would lend itself to routine and standardized steps, in practice the general forensic process for a particular type of crime tends to look similar each time. For example, the examination of digital artifacts to find evidence of child pornography almost always involves taking a bit-for-bit hard drive image and searching common file system locations and slack space for contraband images.

Consider the "Guide to Integrating Forensic Technique into Incident Response" published by NIST (Kent et al. 2006). The NIST process, like many others, can be roughly summarized as follows:

- Collection
- Examination
- Analysis
- Reporting.

Collection involves the process of physical acquisition of data. Examination is the process of combing through the data for items of interest. Analysis is the application of the interesting items to the investigative question at hand, and whether it supports or refutes that question. Reporting describes the output of analysis, including the analysis steps taken.

4 Case Studies

We have developed two hypothetical case studies to reason about the state of digital forensics for cloud-related crimes. While fictional, they describe computer crimes that are not uncommon today. Case Study 1 uses the cloud as an accessory to a crime. Case Study 2 targets the crime against the cloud. These crimes require a reinterpretation when set in a cloud computing environment. In both scenarios, the following themes emerge that differentiate these investigations from traditional digital forensics:

- Acquisition of forensic data is more difficult.
- Cooperation from cloud providers is paramount.
- Current forensic tools appear unsuited to process cloud data.
- Cloud data may lack key forensic metadata.
- Chain of custody is more complex.

We will return to address these issues in more detail in Section 5.

4.1 Case Study 1

Polly is a criminal who traffics in child pornography. He has set up a service in the cloud to store a large collection of contraband images and video. The website allows users to upload and download this content anonymously. He pays for his cloud services with a pre-paid credit card purchased with cash. Polly encrypts his data in cloud storage, and he reverts his virtual webserver to a clean state daily. Law enforcement is tipped off to the website and wishes both to terminate the service and prosecute the criminal.

This is a case where the computer is incidental to the offense. Let us assume that the cloud model used in this case is Infrastructure as a Service, such as Amazon EC2. In this service model, the provider has responsibility and access to only the physical hardware, storage, servers and network components. In the public interest, law enforcement first contacts the cloud provider with a temporary restraining order to suspend the offending service and account, and a preservation letter to preserve evidence pending a warrant.[1] Tracking down the user is the more difficult task. The onus in this case is on the forensic examiner to piece together a circumstantial case based on the data available.

The examiner has no way to image the virtual machine remotely since the cloud provider does not expose that functionality, and in doing so would alter the state of the machine anyway. Deploying a remote forensic agent, such as EnCase Enterprise, would require the suspect's credentials, and functionality of this remote technique within the cloud

---

[1] 18 U.S.C. §2703(f)(1) ("A provider of wire or electronic communication services or a remote computing service, upon the request of a governmental entity, shall take all necessary steps to preserve records and other evidence in its possession pending the issuance of a court order or other process.")

is unknown. Today the forensic examiner, with no case law or standard methodology on the matter, may be tempted to attempt standard practices in digital evidence collection. Namely, with proper recording and documentation, the examiner accesses the offending website and takes snapshots or videotaping the collection of the evidence, and saving the web pages locally. Simply viewing the target website is enough to confirm that the content is illegal, but it tells us nothing about who put it there. Additionally, no guarantee can yet be made that the target webserver has not been compromised by an attacker, or that the examiner's request to the web server was not the victim of DNS poisoning, man-in-the-middle, or some other alteration in transit.

Consider other possible sources of digital evidence in this case: credit card payment information, cloud subscriber information, cloud provider access logs, cloud provider NetFlow logs, the web server virtual machine, and cloud storage data. Law enforcement can issue a search warrant to the cloud provider, which is adequate to compel the provider to provide any of this information that they possess. Law enforcement need not execute or witness the search.[2] The warrant specifies that the data returned be an "exact duplicate," the forensic term that has historically meant a bit-for-bit duplication of a drive. Since child pornography is a federal offense, the provider must comply with the order. A technician at the provider executes the search order from his or her workstation, copying data from the provider's infrastructure and verifying data integrity with hashes of the files. Files may have been distributed across many physical machines, but they are reassembled automatically as the technician accesses them. Though the prosecution may call the technician to testify, we have no implicit guarantees of trust in the technician to collect the complete data, in the cloud infrastructure to produce the true data, nor in the technician's computer or tools used to collect the information correctly. Nonetheless, the provider completes the request, and delivers the data to law enforcement.

Let us say that Polly had two terabytes of stored data.[3] To transfer that quantity of data, the provider saves it to an external hard drive and delivers it to law enforcement by mail. In addition, the provider is able to produce: account information, 10MB of access logs, 100MB of NetFlow records, and a 20GB virtual machine snapshot. After validating the integrity of the data, the forensic examiner is now charged with analysis.

We would expect the forensic expert to identify the following that would aid in prosecution:

- Understand how the web service works, especially how it encrypts/decrypts data from storage
- Find keys to decrypt storage data, and use them to decrypt the data
- Confirm the presence of child pornography
- Analyze logs to identify possible IP addresses of the criminal.

It is not unreasonable to expect that this activity may take many man hours to analyze. According to performance testing from the manufacturer, AccessData found that their Forensic Toolkit (FTK) product took 5.5 hours to process a 120GB hard drive fully on a top-of-the-line workstation, and as long as 38.25 hours on a low-end workstation (AccessData 2010). At that rate, 2TB of data could take 85 hours of processing time. The examiner is likely to dive in first to the data store. The provider may have returned individual files or large files containing "blobs" of binary data. In either case, it will become quickly evident that the data are encrypted. Tools like EnCase and Forensic Toolkit can analyze VMware data files but not snapshots which include suspended memory. The human analyst will have to fix-up and run the VM snapshot in order to understand the website source and observe how encryption is used. Once the keys are uncovered, and data are decrypted, 2TB of data must be analyzed for evidence. We were already aware of illegal content, but not aware of the data owner. Timestamps or file metadata may prove useful, provided they are available and accurate. Evidence of the owner may be gleaned from NetFlow, timestamp, and potentially in the coding style of the website. We can safely assume that an IP can be found that points to Polly. All of the forensic analysis is documented and presented to counsel.

In the absence of legal precedent, existing case law must be considered in the forensic process used. In 2007, the 100-page opinion by Judge Grimm in Lorraine v. Markel issued guidance about the admissibility of original or duplicates of original evidence, as legislated in Rules 1001-1008 of the Federal Rules of Evidence (Lorraine 2007). As mentioned above, service providers are already empowered to conduct searches on behalf of law enforcement. Several important issues regarding the issuance of a warrant were omitted above.

- Search warrants must specify the search of a person or location for evidence of a crime. With cloud computing, a problem emerges because the data may not be location-specific, other than a known public-facing URL or the cloud provider hosting the data. A search warrant must describe the physi-

---

[2] 18 U.S.C. §2703(g)("... the presence of an officer shall not be required for service or execution of a search warrant issued in accordance with this chapter requiring disclosure by a provider of electronic communications service or remote computing service of the contents of communications or records or other information pertaining to a subscriber to or customer of such service.")

[3] Interestingly, 18 U.S.C. §2703(b) allows a cloud provider to disclose the contents of an account used for remote storage without a warrant, and without notifying the customer or subscriber. Kerr suggested that this is unconstitutional (Kerr 2009).

cal place to be searched with particularity.[4] This becomes further complicated if cloud resources are distributed across state or international boundaries.

- The Fourth Amendment presents a preposterous assumption about search preceding seizure,[5] which the courts may be compelled to reinterpret. As Kerr has explored extensively, traditional digital evidence collection is the reverse process of seizure then search (Kerr 2005). Further, digital evidence, and especially cloud evidence, is never "seized" in the sense that it ceases to exist in one place, but the data are the target of the seizure, which are copied and the original remains.

Given the procedure undertaken above, consider the issues which the defense may raise to introduce doubt in the examination:

- Since raw bit-for-bit copies of hard drives were not provided, how do we know that the cloud provider provided a complete and authentic forensic copy of the data? Can the authenticity and integrity of the data be trusted? Can the cloud technician, his/her workstation and tools be verifiably trusted?
- Were the data located on one drive, or distributed over many? Where were the drives containing the data physically located? Who had access to the data, and how was access control enforced? Were the data co-mingled with other users' data?
- If data came from multiple systems, are the timestamps of these systems internally consistent? Can the date and time stamps be trusted, and compared with confidence?
- Does the virtual machine have a static IP address? How can the prosecution tie the malicious activity on the virtual machine to Polly?
- What jurisdiction governs the data in question? If the cloud provider's jurisdiction, then which of their geographic locations or datacenters?

Some of the digital evidence collection from the cloud mirrors traditional collection. In other respects the process is new, such as data dispersed over many storage systems and virtual machine use. Current tools are ill-equipped to process the data in this case easily. The case in almost every

respect hinges on how the cloud provider cooperated. Without greater transparency into how the provider operates, it is difficult or impossible to counter the above objections from the defense.

Finally, we note that cloud providers have a legal obligation to purge child pornography from their systems. Many providers keep duplicate copies of stored data, which here requires that they know where all copies are located and how to verifiably delete the contraband. Microsoft and Amazon declined to comment about their compliance abilities in this situation.

4.2 Case Study 2

Mallory is a hacker who intends to exploit victims by placing a malicious webpage in the cloud. She uses a vulnerability to exploit the cloud presence of Buzz Coffee, a legitimate company. From there, she installs a rootkit that injects a malicious payload into web pages displayed, and hides her malicious activity from the operating system. She then redirects victims to the website, which infects them with malware. Users complain to the legitimate company that they are being infected, so the company seeks to fix the problem and investigate the crime.

This example is a different type of computer crime, one where the target is the computer. Let us assume that Buzz Coffee uses a Software as a Service provider, such as RackSpace. In this service model, the provider has responsibility and access to the hardware, the operating system, and the hosting platform. Buzz wishes to make an example of this hacker, and hires a lawyer to prosecute the attacker. The attorney contracts a forensic specialist to conduct the digital investigation. Using experience as a guide, the investigator constructs a plan to access the cloud provider remotely over a secure channel using Buzz Coffee's credentials and retrieve the website source files. However, when the data are returned, nothing malicious is found since Mallory's rootkit hid the files from the host operating system and the provider's APIs. The forensic investigator determines that the following are additional possible sources of data: cloud provider access logs, cloud provider NetFlow logs, and the web server virtual machine.

The prosecutor approaches the cloud provider with a subpoena and requests all of this data, including a forensic copy of the virtual machine.[6] The provider is willing to conduct an internal investigation; however, it is reluctant to produce the raw data citing confidential and proprietary information. In fact, the Service Level Agreement lacks any language requiring compliance with intrusion response or remediation. The attorney is able to convince a judge that there is

---

[4] Search warrants for online webmail have traditionally specified only the email address as the "place to be searched." See the search warrant for a Gmail mail account at http://docs.justia.com/cases/federal/district-courts/michigan/miedce/2:2009mc50275/237762/2/

[5] "The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no Warrants shall issue, but upon probable cause, supported by Oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized."

[6] Unlike warrants, subpoenas do not require probable cause and can be issued by prosecutors without judicial approval, as long as they are not unreasonably burdensome. See William J. Stuntz, Commentary, *O.J. Simpson, Bill Clinton, and the Transsubstantive Fourth Amendment*, 114 HARV. L. REV. 842, 857-58 (2001).

likely evidence of a crime inside the cloud, and a search warrant is issued to the provider.[7] Even in this case, the provider complies to the extent that its legal counsel feels is appropriate, which in this case includes: NetFlow logs, web access logs, and files from the virtual machine that comprise Buzz Coffee's website. Any further data from the operating system or hosting platform, they claim, would threaten their business and competitive advantage.

A technician at the provider executes the court order from his workstation, copying data from the provider's infrastructure and verifying integrity with MD5 hashes. This information is burned to DVD, and contains 2 MB of NetFlow logs, 100 MB of web access logs and 1 MB of web source code. Using this information, we wish our investigator to uncover the following:

- A chronology that shows when the web pages have been viewed and modified/accessed/created
- Determine the malicious webpage and how the system was compromised
- Analyze the scope of the intrusion, and possible spread to other systems
- Identify the origin of the malicious activity.

Comparing the original website files created by Buzz Coffee to the data returned from the cloud provider would be a constructive first step. Here the technique employed during collection becomes paramount. If the host operating system was used to retrieve the files, Mallory's rootkit would have hidden the malicious files. If files were acquired by reading the physical disk, bypassing the operating system, the complete collection of files will be accurate. Constructing a timeline is a common practice for forensic examiners, and one important in determining when Mallory's files were created. Unfortunately, the procedure employed by the provider again determines whether the investigator receives useful metadata, such as file creation timestamps.

Web access logs are likely the most definitive evidence of the original intrusion, corroborated by NetFlow records. The suspected attacker IP is identified in the logs, which is presented alongside the complete analysis in the subsequent forensic report. Prudent readers might also approach this problem by analyzing the malware installed after visiting the now-hacked webpage, and trying to determine who wrote it or to where it beacons back, but that is not considered here.

Taken to court, the following are questions that could be raised by the defense to discredit the forensic process used in this case:

- Was the chain of custody preserved throughout the process?

---

[7] See examples in NIJ's *Investigations Involving the Internet and Computer Networks*, http://www.ncjrs.gov/pdffiles1/nij/210798.pdf

- Can the malicious page be definitively attributed to Mallory? Who else had access to create/modify this page? Were other clients hosted on the same infrastructure that could have had access?
- What process did the cloud provider use to copy and produce the webpages? Can they make any claims about the forensic integrity of this process? Are timestamps across the different evidence (NetFlow, web logs, etc.) synchronized enough to create an accurate timeline?
- What was the physical location of the virtual machine that is run by the hosting website? By what laws/regulations is it governed?
- What detection and protection mechanisms are employed by the provider to keep their infrastructure secure and to identify intrusions?
- Since the provider refused to provide operating system evidence, can the prosecution have enough evidence to prove that a compromise actually occurred?

In this case the closed nature of the provider was the primary hindrance to a routine investigation. The provider has an incentive to keep as much of its infrastructure private as possible, since it may give them a competitive advantage. Unfortunately, this decision hinders the investigative process and may discredit the legal proceedings that follow.

5 Analysis

Whether in the cloud or not, forensic investigation can be an intensive process. Exams are almost always limited by time and budget, since clients are unwilling or unable to support them indefinitely. Cloud computing, for better or worse, gives customers an ability to terminate virtual machines or revert them to a saved state almost instantaneously. Providers and investigators may also benefit from easy data duplication, system copying/imaging, and extensive business logging. Investigators must recognize the extreme fragility of the evidence. These attributes are indeed positive and contribute towards well-rounded security preparation for incident response. The hindrances seen in these case studies illustrate areas for continued research and development. Consider how we might address the five issues presented at the beginning of Section 4.

First, in our case studies, acquisition was accomplished using legal vehicles of subpoena and search warrant. While somewhat cumbersome given the complex legal system, if a forensic investigation is to support a potential criminal proceeding, this approach is necessary. More efficient mechanisms for the secure transfer of data from providers and law enforcement would be ideal.

Second, cloud consumers will need to negotiate or lobby providers for an appropriate level of cooperation and transparency about how their infrastructure works, the amount of support available during incident response, and forensically-sound practices for assisting law enforcement. One potential

approach is a forensic service level agreement (SLA) appended to the existing SLA signed by providers and subscribers. This legal backing would give customers assurance about the support available to them from their provider during an investigation, a quantitative measure by which to compare providers.

Third, it is clear that remote forensic tools applied to cloud computing are prone to scrutiny, and local processing tools of cloud-stored data are not designed to handle the format or scope of the data. In the case of Infrastructure as a Service, analysis will certainly include the investigation of a virtual machine. Forensic analysts need a tool for parsing, searching and extracting information from virtual machine snapshots, including suspended memory state.

Fourth, the lack of forensic metadata may be addressed in several ways. One proposal is to introduce data provenance in order to track the history and access of cloud objects. In 2007, a report from the Department of Justice recommended asking "what is the chronology of the access to or changes in the data?" of persons providing digital evidence (National Institute of Justice 2007). Another proposal is to introduce preemptive forensics in the cloud, the forensically-sound logging of information at all times without evidence of a crime in order to specifically support forensic investigations after a crime takes place. For example, keeping regular virtual machine snapshots would create a forensic record back in time once an event arises. This computer-generated evidence may benefits from being protected against hearsay arguments, a viewpoint now recognized by some courts.

Finally, chain of custody remains complex given the number of people that may have access to the evidence, and the third-party collection as discussed above. In traditional digital forensics, a chain of custody exists for both physical evidence (*e.g.* the computer) and its associated data. In the cloud case, data are the only evidence. As such, pristine copies of the data, and associated integrity information like MD5 checksums, must be carefully handled. Since chain of custody is the legal equivalent of secure provenance, transfers of custodianship could be documented by a digital provenance system.

Note that we have not addressed the issue of responsibility and fault in either case study. In Case Study 1, we have not established what liability the cloud provider has for hosting the illegal content. In all likelihood, the cloud provider demonstrated no negligence, and is simply a data custodian unaware of the activity. Nonetheless, the law demands they identify and remove all illegal content. In Case Study 2, can users who were infected sue the legitimate company or the cloud provider for negligence? Could Buzz coffee sue the hosting provider if they failed to secure their infrastructure, or to notice the intrusion? These questions may be answerable using an interpretation of current laws. Additionally, we have not explored the investigative complexity of cloud service resellers who themselves offer services that utilize cloud technology. The layering of providers may further complicate the preservation and acquisition of evidence.

Finally, both case studies assume trust in the provider, its employees and infrastructure. Providers have their business reputation and customer base to lose if trust is lost in their ability to provide secure and reliable service. However, if an adversary or corrupt insider gains control over the cloud infrastructure—particularly the hypervisor—no data or computational results in the hosted virtual machines can be trusted.

6 Conclusions

Cloud security is a much discussed topic, but planning about incident response and forensics needs to happen in parallel. The move of data and services to the cloud is already underway, and research and development in the forensic research community must keep pace. These two case studies illustrate larger issues that exist beyond the scope of our specific examples. Forensic acquisition is a renewed challenge, one unsuited for today's tools, which will possibly be addressed by a combination of technological and legal approaches. We have begun to evaluate the ability of popular forensic tools to obtain evidence from a cloud environment. Cooperation with providers will empower consumers to understand their risks and give them leverage to prosecute crimes. The preservation and availability of forensically-relevant metadata remains an open problem.

We have highlighted the issues of common crimes that vary from today only in their use of the cloud. This technology alone introduces peculiarities and open problems that demand immediate attention. As we have shown, deficiencies in both law and technology can be addressed with proper advances.

Acknowledgments

References

AccessData (2010), "FTK Performance Testing," http://www.accessdata.com/downloads/media/FTK Performance Testing.pdf, accessed December 10, 2010.

Bagh, C. (2011), "Amazon EC2 helps researcher to crack Wi-Fi password in 20 minutes," http://www.ibtimes.com/articles/100314/20110112/amazon-ec2-password-wi-hacking-cracking-brute-force-attack-wpa-psk-encryption-cloud-computing-iaa.htm, accessed January 12, 2011.

Choo, K.-K. R. (2009), "Cloud computing: Challenges and future directions," Trends & Issues in Crime and Criminal Justice, no. 400. Canberra, ACT, Australia, October, 2009.

Facebook (2011), "Help Center: How can I download my information from Facebook?" http://www.facebook.com/help/?page=18830, accessed January 4, 2011.

Garfinkel, S. L. (2010), "Digital forensics research: The next 10 years," Proceedings of the Tenth Annual DFRWS Conference, August 2 – 4, 2010, Portland, OR.

Guidance Software (2009), "Facebook Chat Examinations," http://www.encaseondemand.com/EnCast/EnCastVideos/tabid/1383/ProductID/99/CategoryID/129/List/1/Level/1/Default.aspx, accessed January 6, 2011.

Kent, K., Chevalier, S., Grance, T. and Dang, H. (2006), "Guide to Integrating Forensic Techniques into Incident Response," http://csrc.nist.gov/publications/nistpubs/800-86/SP800-86.pdf, accessed September 8, 2010.

Kerr, O. S. (2009), "Applying the Fourth Amendment to the Internet," Stanford Law Review, vol. 62, no. 4, 2009, pp. 1005-1049.

Kerr, O. S. (2005), "Search Warrants in an Era of Digital Evidence," Mississippi Law Journal, vol. 75, 2005, pp. 85-135.

Lemos, R. (2010), "Cloud-Based Denial Of Service Attacks Looming, Researchers Say," http://www.darkreading.com/smb-security/167901073/security/perimeter-security/226500300/index.html, accessed August 4, 2010.

Lillard, T. V. (2010), Digital Forensics for Network, Internet and Cloud Computing: A Forensic Evidence Guide for Moving Targets and Data, Syngress, Rockland, MA.

"Lorraine v. Markel American Insurance Company," 241 F.R.D 534 (D.Md. May 4, 2007).

National Institute of Justice (2007), "Digital Evidence in the Courtroom: A Guide for Law Enforcement & Prosecutors," http://ncjrs.gov/pd_les1/nij/211314.pdf, accessed September 7, 2010.

Ristenpart, T., Tromer, E., Shacham, H. and Savage, S. (2009), "Hey, you, get off my cloud: Exploring information leakage in third-party compute clouds," Proceedings of the 16th ACM Conference on Computer and Communications Security (CCS '09), New York, NY, pp. 199-212.

Schwerha, J. J. and Inch, S. (2008), "Remote Forensics May Bring the Next Sea Change in E-discovery: Are All Networked Computers Now Readily Accessible Under the Revised Federal Rules of Civil Procedure?" Journal of Digital Forensics, Security and Law, vol. 3, no. 3, 2008, pp. 5-28.

Selamat, S., Yusof, R. and Sahib, S. (2008), "Mapping Process of Digital Forensic Investigation Framework," International Journal of Computer Science and Network Security, vol. 8, no. 10, 2008.

Symantec (2011), "The Trojan.Hydraq Incident: Analysis of the Aurora 0-Day Exploit," http://www.symantec.com/connect/blogs/trojanhydraq-incident-analysis-aurora-0-day-exploit, accessed January 21, 2011.

Wolthusen, S. D. (2009), "Overcast: Forensic Discovery in Cloud Environments," Proceedings of the 2009 Fifth International Conference on IT Security Incident Management and IT Forensics (IMF '09), pp. 3-9, September 15-17, 2009, Stuttgart, Germany.