Computer Forensic Analysis in a Virtual Environment

Derek Bem
Ewa Huebner
University of Western Sydney, Australia

Abstract

In this paper we discuss the potential role of virtual environments in the analysis phase of computer forensics investigations. General concepts of virtual environments and software tools are presented and discussed. Further we identify the limitations of virtual environments leading to the conclusion that this method can not be considered to be a replacement for conventional techniques of computer evidence collection and analysis. We propose a new approach where two environments, conventional and virtual, are used independently. Further we demonstrate that this approach can considerably shorten the time of the computer forensics investigation analysis phase and it also allows for better utilisation of less qualified personnel.

Keywords: Computer Forensics, Virtual Machine, computer evidence.

Introduction

In this paper we examine the application of the VMWare (*VMWare*, 2007) virtual environment in the analysis phase of a computer forensics investigation. We show that the environment created by VMWare differs considerably from the original computer system, and because of that VMWare by itself is very unlikely to produce court admissible evidence.

We propose a new approach when two environments, conventional and virtual, are used concurrently and independently. After the images are collected in a forensically sound way, two copies are produced. One copy is protected using the strict chain of custody rules, and the other is given to a technician who works with it in a virtual machine environment not constrained by formal forensics procedures. Any findings are documented and passed to a more qualified person who confirms them in accordance with forensics rules. An additional advantage is that the virtual machine environment makes it easy to demonstrate the findings to a non-technical audience.

An example scenario is described to illustrate our approach. We took a small Windows XP system, created a forensic image of its hard disk, and demonstrated the advantages of using two environments. The example shows that the correct application of a virtual environment approach results in a less time spent on analysing the evidence, giving more chance of discovering important data, and allowing less qualified personnel to be involved in a more productive way. We decided to use only free and readily available utilities to allow everyone to repeat our experiment, and to encourage the reader to try experimenting with their own cases.

What is a Virtual Machine

Virtual machine (also known as 'VM') is a software product which allows the user to create one or more separate environments, each simulating its own set of hardware (CPU, hard disk, memory, network controllers, and other components) and its own software. Ideally each virtual machine should behave like a fully independent computer with its own operating system and its own hardware. The user can control each environment independently and, if required, network virtual computers together or connect them to an external physical network.

While this approach is powerful and flexible, it requires a lot of additional resources, because each virtual computer uses real hardware components present in the computer it runs on. It should also be noted that virtual machine software is complex, and many compromises and restrictions are to be expected. Anyone attempting to use it should have a good understanding of what can and cannot be achieved.

Virtualisation is an old concept, first introduced in the 1960s with the appearance of mainframe computers. It was re-introduced to personal computers in the 1990s, and currently major products available are: Microsoft Virtual PC (*Microsoft Virtual PC 2007*), VMWare software tools range (*VMWare*, 2007), an open source (free) software QEMU (Bellard, 2007), and a few others.

Computer Forensics And Virtual Machine Environments

The conventional computer forensics process comprises a number of steps, and it can be broadly encapsulated in four key phases (Kruse II & Heiser, 2002):

- Access
- Acquire
- Analyse (the focus of this paper)
- Report

During the acquire phase an investigator captures as much live system volatile data as possible, powers down the system, and later creates a forensic (bit by bit) image of all storage devices (Brown, 2005). An image of a storage device is typically acquired using one of many dd based tools (Nelson, Phillips, Enfinger, & Steuart, 2006). This image is stored in the dd format (Rude, 2000), or a proprietary format typically based on dd (Bunting & Wei, 2006). The image is an identical copy of the original disk. It should be noted, however, that the old rule where the image of a hard disk was assumed to be identical with the original hard disk does not necessary apply today. There are many proprietary formats commonly used today which are arguably not identical with the original hard disk; they may include additional metadata like the investigator's name, notes, or hash values. An example of proprietary format is a recently developed and becoming increasingly popular Advanced Forensic Format (AFF) (Garfinkel, 2005). The AFF goes even further by segmenting the original image where each segment has a header, a name, a 32-bit argument, an optional data payload, and finally a tail. The relevance of this short image format overview is the realisation that computer specialist findings may also

be based on examining an image which is in some ways changed, and is not identical with the original.

Because the dd image is the same as the original, it could be copied to the same or a larger hard disk, and booted on another computer system. Such an approach is impractical in recreating the original environment due to too many possible hardware combinations. If the image is booted on a machine with a different hardware configuration, the operating system would discover these differences, and attempt (in some cases unsuccessfully) to install the missing drivers. Furthermore some installed services and software products may refuse to start, or the system could fail to boot at all.

Similar issues exist in a VM environment. VM simulates only some basic hardware components; it is not created to provide full support for a wide range of hardware devices. The acquired dd image can not be immediately booted in a VM environment, as VM requires additional files containing information about the environment being booted. Various software tools can solve this problem by creating the additional files with the parameters required by VM. Some of these utilities are:

- EnCase Physical Disk Emulator (PDE), commercial product (*EnCase Forensic Modules*, 2007),
- ProDiscover family of commercial and free computer security tools from Technology Pathways, LLC (*ProDiscover*, 2007)
- Live View, free utility offered under Gnu Public License (GPL) (*Live View*, 2007)

There were some attempts made to use the VM environment for computer forensics data analysis (ebaca, 2006), but it appears that the suitability of the findings obtained this way as evidence in a court of law is questionable. Some investigators concluded rather prematurely that “VMWare has no real value as a forensic tool” (Fogie, 2004). There are many changes to the original environment required to enable the image to boot in the VM environment, and once the system is booted new data will be written to the original image thus modifying it. An image which is known to be considerably changed would be immediately challenged in a court of law as flawed. A computer expert could argue that the changes were not relevant to the evidence being presented, however it is unlikely that such a line of argument would be accepted by the court. The golden rule “Create a bit-wise copy of the evidence in a backup destination, ensuring that the original data is write-protected. Subsequent data analysis should be performed on this copy and not on the original evidence” is undeniably broken in the virtual environment.

The Proposed Parallel Approach

Each of the four phases of the computer forensics process mentioned in the previous section is further divided into specific steps, and each has to follow strict procedures. The Australian Institute of Criminology guide (McKemmish,1999) recommends that the process of analyzing computer evidence should comply with the following basic rules:

- Minimal handling of the original.
- Account for any change.
- Comply with the rules of evidence.
- Do not exceed your knowledge.

There is no commonly accepted computer forensic certification, but it is expected that a person conducting an analysis and presenting the report in the court is an "expert" (Meyers & Rogers, 2004) who possesses relevant specialised knowledge. We propose that the accuracy of the process can be considerably improved, and the total time required to analyse the data can be shortened if the process is expanded to include two parallel investigative streams, as shown in figure 1.

In our model we used two levels of computer forensics personnel: less experienced and more experienced, respectively referred to as 'Computer Technician' and 'Professional Investigator.' This is similar to the roles CNF Technician (Computer and Network Forensic Technician) and CNF Professional (Computer and Network Forensic Professional) in classification proposed by Yasinsac et al (Yasinsac, Erbacher, Marks, Pollitt, & Sommer, 2003). Modus operandi of such a team is as follows:

- Fully trained and more experienced Professional Investigator adheres strictly to computer forensics investigation methods.
- Less qualified Computer Technician does not have to strictly follow forensic rules, and never has any direct input to the formal reporting process. The role of the Computer Technician is to check the copy of the materials for anything of potential interest and then report the findings to the Professional Investigator.

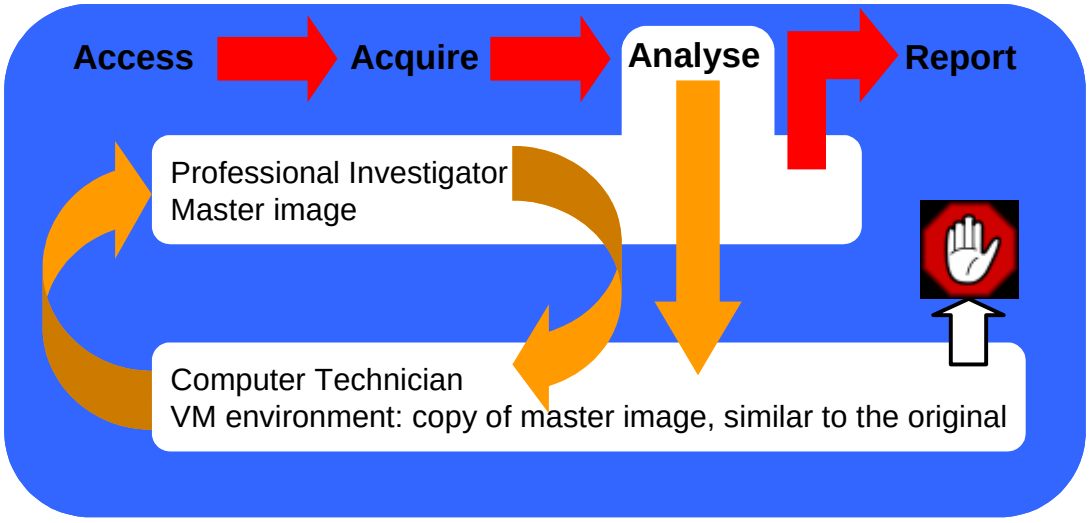

Figure 1: Dual Data Analysis Process

In the analysis phase of the computer forensics investigation a copy of the acquired image is given to the Computer Technician. Their task is to boot the image in a

virtual machine environment, treat it as a normal, 'live' system, and search for all details relevant to the investigation. The methodology used by the Computer Technician invalidates the integrity of the acquired image, but this is of no consequence to the investigation. All findings are passed to the Professional Investigator, who then uses proper computer forensics techniques to confirm the findings, and to make further data searches, if necessary. The findings of the Computer Technician are never included directly in the reporting process. The final report is created by the Professional Investigator, and it only includes the findings confirmed by proper forensic analysis. As we will demonstrate in the following simple example scenario, using the virtual machine environment in tandem with the cooperation between the Computer Technician and the Professional Investigator can deliver better results faster.

The Example Scenario

A powered off personal computer was found when the premises of a person suspected of illegal drug trafficking were searched. A computer forensics investigator was requested to assist in the case. The search warrant provided legal authority, and the investigator was asked to check the computer and find all information pertaining to drug trafficking, including details of financial transactions and any relevant letters or documents. The investigator documented the personal computer's hardware configuration, and acquired the hard disk image using the dd utility from the HELIX bootable forensic CD (E-fense, 2007). SHA-1 and MD5 hash values and relevant case details were recorded, and the chain of custody was created according to local forensic procedures (Hart, April 2004). Two copies of the image named hda1-img3.dd were given to two people in the forensic lab:

- the Professional Investigator, who updated the chain of custody document and locked the image in a safe place,
- the Computer Technician, who updated the chain of custody corresponding to the second image, and also locked the image in a safe place.

The Computer Technician was asked to boot the image in a VM environment, and to use any suitable tools to search the booted system. To facilitate this VMWare Server (*VMWare Server*, 2007), a free virtualisation product from VMWare, was installed on a separate computer. As expected, the image hda1-img3.dd could not be directly booted in the VM environment. The Computer Technician used Live View software, see figure 2, to create additional files needed to boot the image in VM. The Live View proved to be simple to use and reliable, but any other software with the same functionality could also be used. The Live View messages window creates a comprehensive activity log, which was copied by the Computer Technician and included in the notes. After Live View finished creating files, it automatically offered to boot the image in VMWare.

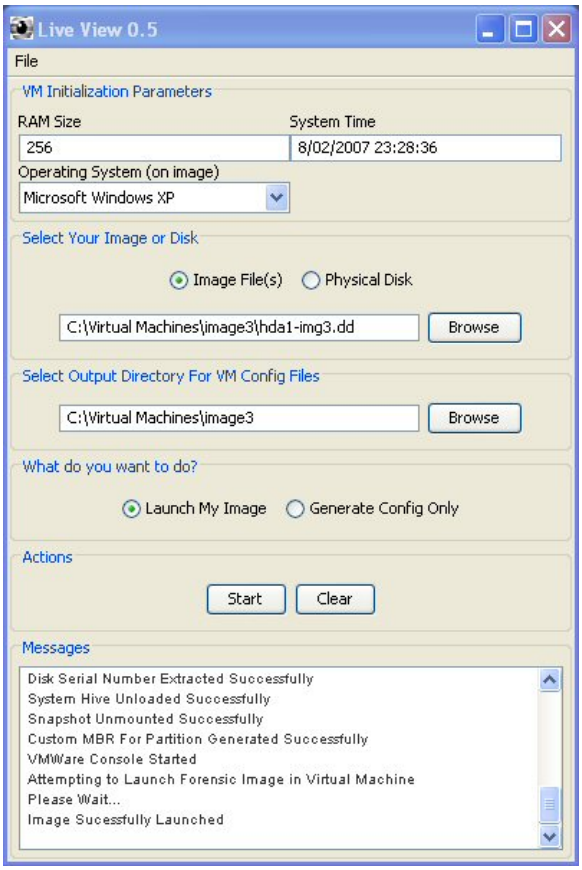

Figure 2: Live View Window

The first boot of hda1-img3.dd image in the VMWare Server produced a series of 'found new hardware' messages. This was expected, as VMWare simulates different hardware than the hardware of the original Windows XP installation. New devices were identified and installed; all required drivers were common, as they were a part of the standard Windows distribution, and were quickly installed without the need to provide propriety software.

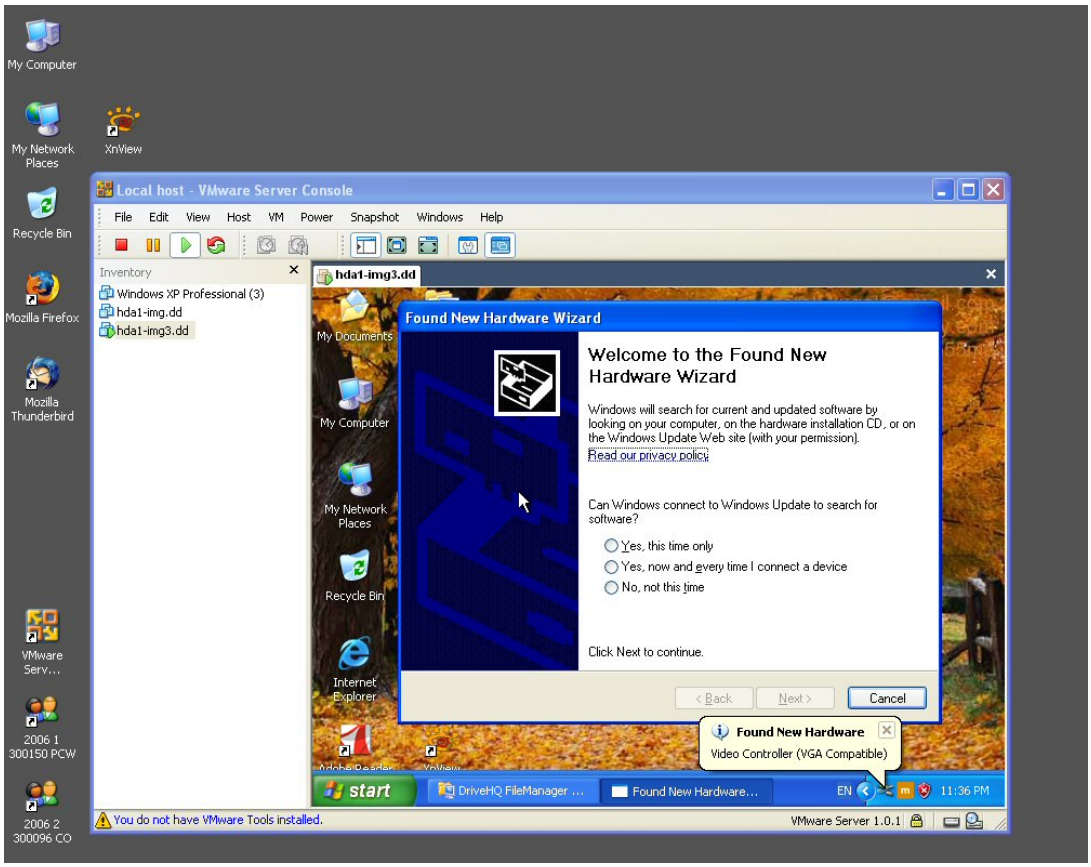

Figure 3: VMWare Server Window

After the first successful boot of the system the Computer Technician powered the VM machine off, and installed VMWare tools which improve mouse operation and provide a higher VM screen resolution (see the message “You do not have VMWare Tools installed” in left hand bottom corner in figure 3). Before booting the system again the Computer Technician checked the virtual machine settings shown in figure 4, and noted that only four devices were installed: memory, hard disk, CD-ROM and the USB controller. Other devices available to the virtual machine were seen in the Add Hardware Wizard, but they were of no immediate interest. The Ethernet controller was not installed, thus the virtual machine was isolated from any networks and the Internet.

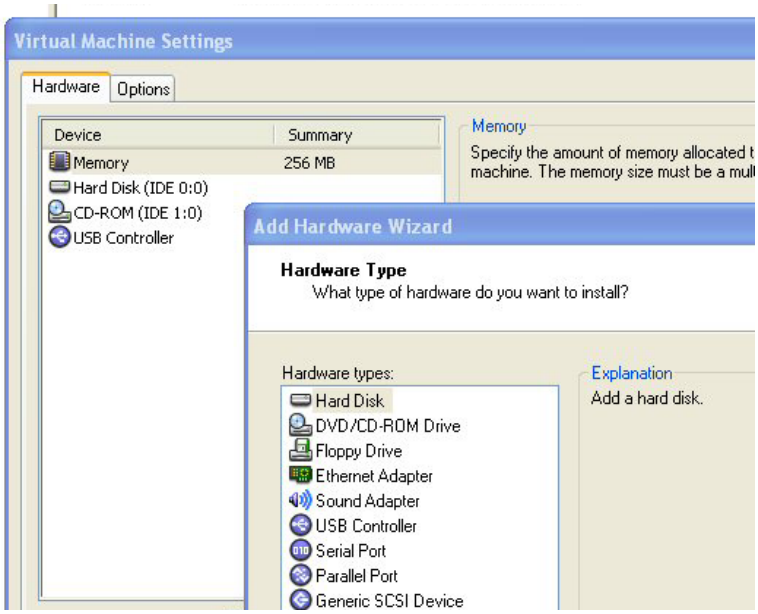

Figure 4: Virtual Machine Settings

Each time the system was booted the Computer Technician observed the panel which started automatically and indicated an attempt to log on to DriveHQ, as shown in figure 5. These attempts failed, as there was no Internet connection from this virtual machine.

Figure 5: DriveHQ Log On Window

The Computer Technician checked the Internet from another computer, and found that the DriveHQ ("Drive Headquarters", 2007) is an Internet virtual storage service which allows a user to store a large volume of any data. To access the DriveHQ account a user name and password were required. The user name ‘kugel’ was visible in the log on panel, but the password was hidden behind a row of dots. In order to find out what the password was the Computer Technician decided to install in the virtual machine additional software called ‘password revealer.’ There are many password revealing tools, and they have different capabilities; in this case some of them failed to reveal the password. Finally a tool named Aqua Deskperience succeeded (*Aqua Deskperience*, 2007): the password was “just555me” (see figure 6).

Figure 6: Aqua Deskperience Reveals Drivehq Account Password

The Computer Technician continued the examination of the booted system using standard Windows tools. Windows Explorer did not show any folders or files of relevance to the case being investigated. Web browser bookmarks also did not point to any Web sites which could be relevant.

The Computer Technician then checked what additional software was installed on the investigated system, and discovered on the desktop an icon for the Simple File Shredder (*Simple File Shredder*, 2006) (see figure 7). A quick Internet check showed that the Simple File Shredder is a utility that securely deletes files, making their recovery impossible.

Figure 7: Desktop icons

Two important points are worth noting here:

1. The image of the disk used to run the system is no longer identical with the image hda1-img3.dd acquired using forensically sound methods. We installed various new device drivers and new software packages (Aqua Deskperience, possibly a few others as well). Various files and folders were 'touched' by checking them in Windows Explorer, and by opening them in their native applications. It would be unrealistic to argue that the image is still valid as evidence; it is now contaminated.
2. The original acquired image hda1-img3.dd is still kept in custody by the Professional Investigator; it is unchanged and forensically valid.

The Computer Technician reported to the Professional Investigator the following basic findings:

- It is likely that there are not many traces of any files of interest to the investigation, as the person using the computer installed a secure deletion tool. It is possible that all such files have been irrecoverably erased.
- It is likely that materials of evidentiary value were not kept on the computer, as the owner had a virtual Internet storage account which allowed them to keep all data on a remote server.
- It was possible to recover the user name and password to the remote storage system account.

Using this information the Professional Investigator can now confirm all the findings using forensically sound methodology and proper forensics software tools. The remote account on DriveHQ can be accessed from another computer, and all files stored there copied and examined. The hda1-img3.dd disk image can still be analysed with computer forensics software for any traces of unshredded files with the knowledge that the likelihood of finding anything of value is small and a large amounts of time should not be dedicated to this task.

Conclusion

The simple scenario presented above demonstrates that the cooperation between two teams equipped with different sets of tools, and using personnel with different levels of expertise, can produce much faster results, and will lessen the workload of highly qualified Professional Investigators. The Professional Investigator using proper computer forensics tools and techniques would most likely achieve the same results working in a conventional setup, without using a virtual environment and without the help of a Computer Technician. However the described method of using two environments, conventional and virtual, could save time and increase the chances of finding important evidence.

If only conventional image analysing techniques were used in the presented example scenario, considerable skills and significantly more time would be required to find the DriveHQ account, the user name and password. The same information was in plain view when the Computer Technician booted the investigated image in VMWare, and the technician used commonly available tools (e.g. the password revealer) to search the investigated system. It was then considerably easier for the Professional Investigator to benefit from the initial findings, and to conduct a forensically sound and properly documented search on the image.

An additional advantage for an organization is that technical personnel can be exposed to computer forensics techniques in stages, without compromising real evidence, yet at the same time providing genuinely valuable input to the process. This approach can be also seen as a part of an internal training process, where a person with little computer forensics experience, but good technical knowledge, is not immediately given the responsibilities of a Professional Investigator, but is first introduced to the forensic process by conducting investigations in a virtual environment.

In this paper we described the process of using conventional and virtual environments in the analysis phase of computer forensics investigations. We also proposed the ground rules for cooperation between the Computer Technician and the Professional Investigator. We believe that future research is needed to better formalise the whole process, with emphasis on what is expected from the Computer Technician. Future research in the area is also required to more thoroughly test other available virtualisation software tools, and to find their strengths and weaknesses.

About the Authors

Derek Bem, MElecEng Warsaw, MIEAust, CPEng (d.bem@scm.uws.edu.au) is a Lecturer in the School of Computing and Mathematics at University of Western Sydney, Australia. Derek has extensive experience in the computer industry, where he worked in IBM and other companies as a hardware and software engineer. Derek is currently coordinating UWS teaching in the Computer Forensics area.

Ewa Huebner, MElecEng Warsaw, PhD Sydney, MACS (e.huebner@scm.uws.edu.au) is a Senior Lecturer in the School of Computing and Mathematics at University of Western Sydney, Australia. Ewa has extensive experience teaching computer science subjects, and she is currently leading the UWS Computer Forensics research group. UWS computer forensics web site can be found at http://www.scm.uws.edu.au/compsci/computerforensics/.

References

Deskperience. (2007). *Aqua Deskperience*. Retrieved 20 January 2007, from http://www.deskperience.com/aqua/index.html

Bellard, Fabrice. (2007). *QEMU*. Retrieved January 17, 2007, from http://fabrice.bellard.free.fr/qemu/index.html

Brown, C. L. T. (2005). *Computer Evidence: Collection & Preservation*. Hingham, MA: Charles River Media.

Bunting, S., & Wei, W. (2006). *EnCase Computer Forensics: The Official EnCE: EnCase Certified Examiner Study Guide* (1st ed.). Indianapolis, IN: Wiley Publishing.

CERT, Software Engineering Institute. (2007). *Live View*. Retrieved February 12, 2007, from http://liveview.sourceforge.net/

Drive Headquarters. (2007). Retrieved 2 November 2006, 2 November 2006, from http://www.drivehq.com/

ebaca. (2006). *Penguin Sleuth Kit Virtual Computer Forensics and Security Platform*. Retrieved 28 November 2006 from http://www.vmware.com/vmtn/appliances/directory/249

E-fense. (2007). The HELIX Live CD Page. Retrieved 9 February 2007, from http://www.e-fense.com/helix/

Fogie, S. (2004). *VOOM vs The Virus (CIH)*. Retrieved 12 March 2005 from http://www.voomtech.com/VOOM_vs_The_Virus.html

Garfinkel, S. (2005). *The Advanced Forensic Format 1.0*. Retrieved November 13, 2006 from http://www.afflib.org/affdoc.pdf

Guidance Software. (2007). *EnCase Forensic Modules*. Retrieved January 25, 2007, from http://www.guidancesoftware.com/products/ef_modules.asp

Hart, S. V. (April 2004). *Forensic Examination of Digital Evidence: A Guide for Law Enforcement*. from http://www.ojp.usdoj.gov/nij/pubs-sum/199408.htm

Kruse II, W. G., & Heiser, J. G. (2002). *Computer Forensics: Incident Response Essentials* (1st ed.): Addison Wesley Professional.

McKemmish, R. (1999). *What is Forensic Computing?* : Australian Institute of Criminology.

Meyers, M., & Rogers, M. (2004). Computer Forensics: The Need for Standardization and Certification. *International Journal of Digital Evidence, 3*(2).

Microsoft. *Microsoft Virtual PC 2007*. Retrieved 22 February 2007, from http://www.microsoft.com/windows/products/winfamily/virtualpc/default.mspx

Nelson, B., Phillips, A., Enfinger, F., & Steuart, C. (2006). *Guide to Computer Forensics and Investigations, Second Edition* (2nd ed.). Boston, MA: Thomson Course Technology.

Rude, T. (2000). *DD and Computer Forensics*. Retrieved October 23, 2003 from http://www.crazytrain.com/dd.html

scar5 Software. (2006). *Simple File Shredder*. Retrieved 15 December 2006, from http://www.scar5.com/

Technology Pathways, LLC. (2007). *ProDiscover*. Retrieved January 2, 2006, from http://www.techpathways.com/

VMWare. (2007). *VMWare*. Retrieved February 14, 2006, from http://www.vmware.com/

VMWare. (2007). *VMWare Server*. Retrieved February 15, 2006, from http://www.vmware.com/products/server/

Yasinsac, A., Erbacher, R. F., Marks, D. G., Pollitt, M. M., & Sommer, P. M. (2003). Computer Forensics Education. *IEEE Security and Privacy, 1*(4), pp. 15-23.