Trends in Virtualized User Environments

Diane Barrett
University of Advancing Technology
Dbarrett@uat.edu

ABSTRACT

Virtualized environments can make forensics investigation more difficult. Technological advances in virtualization tools essentially make removable media a PC that can be carried around in a pocket or around a neck. Running operating systems and applications this way leaves very little trace on the host system. This paper will explore all the newest methods for virtualized environments and the implications they have on the world of forensics. It will begin by describing and differentiating between software and hardware virtualization. It will then move on to explain the various methods used for server and desktop virtualization. Next, it will explain how virtualization affects the basic forensic process. Finally, it will describe the common methods to find virtualization artifacts and identify virtual activities that affect the examination process of certain virtualized user environments.

Keywords: Hardware-assisted, Hypervisor, Para-virtualization, Virtual Machine, virtualization, VMware, Moka5, MojoPac, Portable Virtual Privacy Machine, VirtualBox.

1. INTRODUCTION

According to a research published by Gartner in February of this year, there are nearly 100 providers of products adapted for the server virtualization management marketplace [1]. Fewer than 5 million PCs were "virtualized" in 2006; by 2011, that figure will rise to between 480 million and 846 million [2].

With more emphasis being placed on going green and power becoming more expensive, virtualization offers cost benefits by decreasing the number of physical machines required within an environment. A virtualized environment offers reduced support by making testing and maintenance easier. On the client side, the ability to run multiple operating environments allows a machine to support applications and services for an operating environment other than the primary environment. This decreases costly upgrade costs and allows more uniformity in desktop environments.

In late 2007, the Distributed Management Task Force, Inc. created an open standard for system virtualization management. This standard recognizes supported virtualization management capabilities for discovering virtual computer systems, managing the lifecycle of virtual computer systems, controlling virtual resources and monitoring virtual systems. These developments should be of interest to the digital forensic investigator for several reasons. An increase in the use of virtual environments and applications that can be run from a USB device means that any incriminating evidence may not be readily found, especially if the device itself is not recovered. There is an increased interest in the use and detection of virtual machine environments (VMEs) by those that want to spread malware or conceal activities. When malicious code is released that makes use of its own VME, it will become essential for anti-malware researchers to find ways to detect the VME. Additionally, computer forensics professionals will be required to detect and examine such environments.

2. HOW VIRTUALIZATION WORKS

In order for virtualization to happen, a hypervisor is used. The hypervisor controls how access to a computer's processors and memory is shared. A hypervisor or virtual machine monitor (VMM) is a virtualization platform that provides more than one operating systems to run on a host computer at the same time. This section will take a brief look at the underlying technologies of virtualization.

2.1 Hardware

A Type 1 native or bare-metal hypervisor is software that runs directly on a hardware platform. The guest operating system runs at the second level above the hardware. These hardware-bound virtual machine emulators rely on the real, underlying CPU to execute non-sensitive instructions at native speed [3]. In hardware virtualization, a guest operating system is run under control of a host system, where the guest has been ported to a virtual architecture which is almost like the hardware it is actually running on. The guest OS is not aware it is being virtualized and requires no modification. The hypervisor translates all operating system instructions on the fly and caches the results for future use, while user level instructions run unmodified at native speed [4].

2.2 Paravirtualization and Hardware Assist

Paravirtualization involves modifying the OS kernel to replace nonvirtualizable instructions with hypercalls that communicate directly with the virtualization layer hypervisor. The hypervisor also provides hypercall interfaces for other critical kernel operations such as memory management and interrupt handling [5].

The Virtual Machine Interface (or VMI) was developed by VMware as a mechanism for providing transparent paravirtualization. The VMI interface works by isolating any operations which may require hypervisor intervention into a special set of function calls. The implementation of those functions loads a "hypervisor ROM" [6]. This design also allows the same binary kernel image to run under a variety of hypervisors, or, with the right ROM, in native mode on the bare hardware.

Vendors are currently working on second generation hardware assist technologies that will have a greater impact on virtualization performance while reducing memory overhead.

2.3 Hosted hypervisors

A Type 2 or hosted hypervisor is software that runs within an operating system environment and the guest operating system runs at the third level above the hardware. The hypervisor runs as an application or shell on another already running operating system. Operating systems running on the hypervisor are then called guest or virtual operating systems. This type of virtual machine is composed entirely of software and contains no hardware components whatsoever. Thus, the host can boot to completion, and launch any number of applications as usual, with one them being the virtual machine emulator.

3. VIRTUAL TECHNOLOGY IN BUSINESS

As more and more vendors venture into the realm of virtualization, forensic investigators will be faced with not only doing network forensics, but doing virtual network forensics as well. IDC Research predicts that spending on virtualization will reach almost $15 billion worldwide by 2009. Companies are already pushing out virtual desktops made from virtual image files. For example, Cincinnati Bell decided that desktop virtualization is a better alternative to upgrading hundreds of PCs running Windows 2000. They can now start up 800 virtual desktops each day using only 12 virtual images.

The virtualization market consists of many products and vendors. The next section lists some of the major players in the market along with the type of virtualization technology they offer.

3.1 VMware and VMware Preconfigured Appliances

VMware offers a wide variety of services and products. Of particular interest are the preconfigured appliances that can be readily downloaded form VMware’s virtual marketplace and implemented using VMware Player. A virtual appliance is a pre-built, pre-configured and ready-to-use enterprise software application on a virtual machine.

3.2 Microsoft Virtual Products

Microsoft’s solution includes servers, desktops, and applications virtual machine management and virtualization acceleration. Recently Microsoft has released it’s own hypervisor called Hyper-V and has begun offering pre-configured Virtual Hard Disks (VHDs) that can be downloaded similar the virtual appliance market of VMware.

3.3 XenSource

XenSource distributes its hypervisor as free, open-source software. XenSource has been acquired by Citrix. It offers a comprehensive end-to-end virtualization solution with the main purpose to enable IT to deliver applications to users anywhere.

3.4 Parallels Preconfigured Appliances

Parallels is similar in VMware in that it provides virtualization solutions along with preconfigured appliances except for the Macintosh platform. The company offers a library with more than 350 software downloads that can be used to create and manage operating systems and applications running in virtual environments.

3.5 Virtualization Boxes

There is a wide variety of other companies that offer virtual solutions. Since the market is growing at a very fast pace, included in this section are only the solutions the author found currently relevant to the computer forensics realm.

Sun Microsystems VirtualBox is a mature virtualization tool that runs on Windows, Linux, Macintosh and Solaris. It supports Windows (including Vista), Linux, OS/2 Warp, OpenBSD, and FreeBSD as guest operating systems.

The Pano device by Pano Logic has no CPU, no memory, operating system, drivers, software or moving parts. It is merely a box that has connections for a keyboard, mouse, display, audio and USB peripherals. It connects over an existing IP network to an instance of Windows XP or Vista running on a virtualized server.

The InBoxer Anti-Risk Appliance is used for email archiving, electronic discovery, and real-time content monitoring. The InBoxer virtual appliance can be run on either VMware server or Microsoft Virtual PC as well.

4. VIRTUAL TECHNOLOGY FOR INDIVIDUAL USE

The use of virtualization is growing in the individual use market as well as the corporate environment. This section explores the technology being used with personal computer that do not alter the current environment, but use a USB device to run the virtual environment, thereby leaving the original system intact.

4.1 MojoPac

MojoPac is developed by RingCube. It can be used as an individual or an enterprise solution. MojoPac’s virtualization technology encapsulates a complete Windows XP desktop environment, isolating it from the underlying host PC. This virtualized environment can be loaded onto a host computer, a portable USB storage device, or network attached storage and run on any Windows host computer.

4.2 Moka5

A Moka5 LivePC contains everything needed to run a virtual computer. LivePCs can be run from a USB flash drive, iPod, or a desktop computer. LivePCs can be downloaded from a repository of LivePCs similar to VMware and Parallels concept. The Moka5 engine streams and prefetchs these files so they can be shared. It also automatically updates the LivePCs as changes are made to them.

4.3 Portable Virtual Privacy Machine

The Portable Privacy Machine by MetroPipe contains a complete virtual Linux machine with privacy-enabled Open Source Internet applications. The Portable Privacy Machine is based on Damn Small Linux (DSL) and QEMU releases [22]. QEMU is a generic, open source processor emulator. As with other products in this category, it can be loaded on USB drives, Flash Memory cards, Secure Digital devices, or iPods.

4.4 Preconfigured virtual appliances

VMware hosts about 725 virtual appliances that can easily downloaded and installed. Microsoft virtual appliances are surfacing and earlier it was mentioned that Parallels offers more than 350 virtual appliances. Available ready to go, are over1,000 virtual appliances that anyone can use.

5. EXAMINING VIRTUALIZED USER ENVIRONMENTS

Traditionally virtual machines have been used to create contained environments for malware isolation or to examine suspect machines. This allows the forensic examiner to boot the image or disk and gain an interactive, user-level perspective of the environment without modifying the underlying image or disk. However, now instead of using virtual environments to examine machines, virtual environments themselves need to be examined.

The research conducted includes exploring what remnants are left by virtual environments that were run from a UBS drive. The environments examined were MojoPac, Moka5, Portable Virtual Privacy machine, and a VMware appliance.

The methodology used was kept as simple as possible to gain an accurate picture of the environments. FTK Imager was used to make a dd image from a clean Windows XP install on a 20GB drive. The USB device was plugged in, the virtual environment was started, several actions were performed such as Internet surfing and then the device was ejected. FTK was used the take another dd image of the machine. The dd image chunks were reconstructed into one file using A.F.7 Merge, and Beyond Compare was used initially to look for differences in the dd images. FTK was also was used to search for signs of the virtual environment. The following sections describe the observations made during examination of these environments.

5.1 MojoPac

Notable findings:

- The NTUSER.DAT file contained the line: Autorun Action Run MojoPac.
- The Windows SysEvent.evt log file contained the phrase: To see your password hint, please move the mouse over the question mark in the MojoPac Login Dialog.
- Three prefetch files that all listed: \DEVICE\HARD DISK1\DP(1)0-0+9\PROGRAM FILES\RINGTHREE\BIN\MOJOPAC.DLL.
- Pvm.sys, ringthree.ico were found stored on the host machine
- Phones home for updates

MojoPac allows for all documents and personal settings to be copied to the drive, before launching. If this happens, there will be .lnk files. Although the application does not allow access to the local hard drive once the application is started, access to the CD/DVD drive and removable drives is still possible. MojoPac implements paging between memory and the hard drive to take place on the host PC instead of on the portable drive, so remnants of activity from the drive would be in the pagefile. Browsing and multimedia history stays inside MojoPac. It has a separate registry and shell stored on the USB device. Currently it will only run on Windows XP and needs administrative rights on the host machine in order to run, unless a application such as MojoUsher is installed on the host PC for limited mode authority. MojoPac runs under the RingThreeMainWin32 process. Since there are essentially 2 XP environments running programs of the same name may be running on both the host and the virtual environments.

5.2 Moka5

Notable findings:

- Creates folders in the My Documents folder for Live PCs and Live PC Documents. These folders are not removed when the drive is ejected.
- Entry in the user’s Startup folder for Moka5 USB Clean 2238, which points to an executable file in the host machine’s C: drive: C:\Documents and Settings\Local Settings\Temp\m5usb-2238\m5usb.exe.
- Folder labeled m5usb-2238 inside the Temp folder which contained a total of 23 Moka5-related files.
- Evidence of registry keys created or modified by Moka5
- Log file containing information on a Moka5 automatic updates client

on the host machine and the path of the Moka5 engine on the thumb drive from which it was run.

- Phones home for updates upon launch
- Live PCs are stored with .lpc extension

Moka5 technology is based on VMware Player. The application asks whether you want to leave it installed for easier load next time, so there will be evidence in both the Temp folder and the Application Data folder with VMware references. The Moka5 Engine will stream and prefetch LivePCs. Any changes made during a session are captured in separate file systems on a ramdisk. Browsing and multimedia history stays inside the virtual machine.

5.3 Portable Virtual Privacy Machine

Notable findings:

- NTUSER.DAT and NTUSER.LOG files changed
- Prefetch data files are present
- Phones home

Portable Virtual Privacy Machine technology is designed to just plug the drive into any Windows or Linux computer, and the Virtual Privacy Machine will run a contained environment including portable applications. All Browsing and multimedia history stays inside the virtual machine. According to a notice posted on MetroPipe’s website, as of July 2008, the version of the Portable Privacy Machine that was tested is no longer maintained nor supported. A new version is in development using updated software and operating system. This environment will be reexamined, once the new version is released.

5.4 Preconfigured VMware virtual appliances

Notable findings:

- runs via VMware Player
- creates two VMware network adapters

When VMware player is used, there will be traces associated with VMware, such as c:\program files\common files\vmware. This type of environment is perhaps the most obvious to spot. There will be virtual adapters created and host of VMware referenced files. The user activity is contained in the virtual appliance.

6. WHAT TO LOOK FOR

In many of the aforementioned technologies, virtual devices are exclusive to the virtual machine and are files on the host. For example, VMware creates virtual adapters as well as files with extensions: .vmx, .vmdk, .vmsn and vmss. “What Files Make Up a Virtual Machine?” posted on VMware’s website is an excellent resource for files extensions that are associated with VMware along with the purpose of the file. Since some forensic software lists these extensions as unknown file types, a forensic examiner should become familiar with these files. Otherwise they can easily be skipped over in an investigation. The same goes for other virtual formats. As the number of vendors that create virtual solutions increases so does the types of image storage formats.

The host's critical resources such as memory, processor time, video, and sound are shared with the virtual machines. In applications such as MojoPac, the host resources must be utilized for better performance. Log files are created by most software; virtual machines are no exception, look for these. Since many of these technologies use a USB drive for access, there will be remnants in the registry. In the March 2007 edition of Digital Investigation an article titled "Tackling the U3 trend with computer forensics" by Andy Spruill and Chris Pavan explores the artifacts left behind by U3 devices. The information provided is a good base for some general items to be on the lookout for:

- MRU cache
- Link files
- Prefetch files
- Page file
- Unique identifiers associated with the program
- Artifacts in processes, file system, and/or registry
- Artifacts in memory
- VME-specific virtual hardware, processor instructions and capabilities

Research conducted found that this list can be used as a starting point. Since individual environments vary, not all these will exist, especially with applications such as MojPac and Virtual Privacy Machine.

In the corporate environment, Application-layer security, such as application proxies can capture some evidence that can help track actions. Application-layer firewall logging can capture more than the IP address and port number. Many firewalls are capable of intercepting packets traveling to or from an application such as a browser. This provides a more thorough examination of network traffic and can capture evidence from applications such as Moka5 and Portable Virtual Privacy Machine. Corporations also have the option of not allowing removable media. This can eliminate the issues that arise from using many of the technologies mentioned here.

The home environment becomes a bit more difficult. If the user is computer savvy, finding tracks may be almost impossible. Devices are becoming smaller with larger capacity and can easily be hidden. Home environments need to be examined very closely for all CDs and removable devices.

7. CURRENT CHALLENGES

A virtual machine located inside a forensic image cannot be properly examined by most software. Forensic software reports the virtual machine files as unknown file types. Although the virtual machine can be exported or loaded into another virtual machine, when that suspect virtual machine is loaded the file information inside the original virtual machine changes.

7.1 File format conversion

EnCase allows a .vmdk file to be added as an evidence file for analysis. In order to do this, the .vmdk file can be export out and then add back in separately. Once the .vmdk file is added into the case, EnCase sees it as the hard drive of the virtual machine. The research into the examination of these various environments included the quest for programs that would convert a virtual image to a more universal format such as a dd file. This was done to find a way to convert the environment for mounting and examination without changing the original files. FTK Imager will open .vmdk files and acquire it to dd image. In addition to forensic software, programs such as Live View can mount the image write protected so that no alterations are done to that DD image. Though, this is a start, not all virtual environments are this easy to examine.

The experimentation process for examining a .vhd file used WinImage to convert and mount the virtual machine file. The following describes the steps taken:

- Retrieve virtual disk image from target machine
- Hash the image for access control, chain of custody
- Access disk with WinImage, extract files as necessary
- Hash the disk image a second time to verify that WinImage did not modify the original virtual disk
- Load extracted files into FTK or forensic tool of choice for analysis

This process did not modify any of the files as the hash values matched. Due to time constraints, other file formats were not tested. This research is ongoing. The challenge is finding utilities that recognize and convert file formats.

7.2 Recognizing virtual environments

In his presentation on the Effectiveness of Hash Sets, Douglas White of the National Institute of Standards & Technology (NIST) compares physical and virtual OS installations. There is a difference in the number of files in each type of installation. His research shows the differences in physical vs. virtual machines appear to be due to virtual machines using abstract or generic device interfaces and physical machines requiring vendor specific drivers.

This being said, any investigation now must first determine if the device being examined is real or virtual. In dead drive forensics the virtual machine file itself will be present. In live forensics, the differences may be a bit more subtle as the virtual environment may be running. Determining if the environment is real can be done in several ways. In November of 2004, Joanna Rutkowski published the Red Pill or how to detect VMM using (almost) one CPU instruction [7]. The Red Pill focuses on detecting virtual machine usage without looking for file system artifacts. It is based on relocation of sensitive data structures. Scoopy Doo and Jerry are tools that detect a VMware fingerprint. When Scoopy Doo is run, it simply states: This is/is not a virtual machine. These tools can be found at: http://www.trapkit.de/research/vmm/index.html. On this website, Tobias Klein also poses the question “is it possible to break out of a VM (to reach the Host OS or to manipulate other VMs)? This is quite an interesting question as the implications can be great since virtualization is based on isolated environments. This has become a recent topic of discussion on security forums. Articles about the security of virtualized environments are beginning to surface. For those more adventurous, Snoopy Pro is available. This tool analyzes virtual traffic between the device and driver.

When examining virtualized environments, it is important to reflect on what is being captured. Tools available to examine virtual environments are limited. The Volatility Framework 1.1.1 is a collection of tools, for the extraction of digital artifacts from volatile memory (RAM) images. The framework is intended to introduce people to the techniques and complexities associated with extracting digital artifacts from volatile memory images and provide a platform for further research into this area [8]. In May 2007 Network General added virtual server forensics. The company added modules that let IT personnel peer into the workings of VMware's ESX and Microsoft Virtual Server. However, when Henderson and Dvorak, who are members of the Network World Lab Alliance, tested the virtual-machine-monitoring capabilities, they found it takes a lot of preparation and configuration work to yield useful data [9].

8. WHAT THE FUTURE HOLDS

Virtualization appears to have a definite hold on the market and companies are competing fiercely to develop and implement products for this environment. Along with all these changes and technologies, challenges will come.

Our court system already has a difficult time with cyber crime. Earlier this year a federal grand jury issued a subpoena to MySpace.com in a case where a teenage girl committed suicide. Federal prosecutors are charged the suspect with defrauding MySpace for creating a false account in CA, because the state where the crime happened did not have legal violations that could be prosecuted. In another recent cyber crime case, the judge ruled that there was no crime because it was a faceless crime. The judicial system is hard pressed to keep up with the changing face of digital crime. The possibility of a challenge based on virtual environments exists.

What happens when we have kiosks that a user downloads a virtual environment into a browser, commits a crime, and then deletes the virtual machine? This can happen anywhere. Virtual social networks continue to grow. How will crime be investigated in Second Life? There are accounts of money laundering, identity stealth and intellectual property theft that result from actions that occur in a virtual world.

In virtualization, there is the ability to roll back or delete a bad or defective machine. With the Federal Rules of Civil Procedure governing data retention, will virtual machines need to be included in an organizations data retention policy? We are moving to more of a dynamic environment where organizations are pushing out virtual operating systems to desktops from a server and streaming applications on an as needed basis. At the end of day everything goes away and the user environment starts fresh the next day. Microsoft finds the idea that you can make pools of dynamic resources with unlimited capacity available to users anywhere at any time extraordinary. For an investigator, this environment can be quite complicated and a bit unnerving.

As investigators find ways to examine virtual machines, will the processes be questioned as to the original evidence file? Borrowing the last line from "Attacks on More Virtual Machine Emulators" by Peter Ferrie: "One thing is clear – the future looks complicated".

REFERENCES

[1] Gartner Research, The Server Virtualization Management Marketplace. Publication Date: 19 February 2008, ID Number: G00154109.

[2] Gammage, B., Shiffler III, G. Report Highlight for Dataquest Insight: PC Virtualization Forecast Scenarios. Gartner Research, Publication Date: 8 August 2007 ID Number: G00150832.

[3] Ferrie, P. n.d. Attacks on More Virtual Machine Emulators. www.symantec.com/avcenter/reference/Virtual_Machine_Threats.pdf.

[4] Paravirtualization API Version 2.5. Copyright 2005, 2006, VMware, Inc. www.vmware.com/pdf/vmi_specs.pdf.

[5] Understanding Full Virtualization Paravirtualization and Hardware Assist. www.vmware.com/files/pdf/VMware_paravirtualization.pdf.

[6] The VMI virtualization interface. http://lwn.net/Articles/175706/. Posted March 15, 2006 by corbet.

[7] Rutkowski, J. Red Pill... or how to detect VMM using (almost) one CPU

instruction http://invisiblethings.org/papers/redpill.html.

[8] Volatility Framework 1.1.1 (GPL). http://www.nabble.com/Volatility-Framework-1.1.1-(GPL)-td12136727.html.

[9] Henderson, T, Dvorak, R. Network General tool peers inside virtual machines, Network World, 07/09/07, http://www.networkworld.com/reviews/2007/070907-network-general-test.html?page=1.