

available at www.sciencedirect.comjournal homepage: www.elsevier.com/locate/diinDigital
Investigation

A survey of main memory acquisition and analysis techniques for the windows operating system

Stefan Vömel*, Felix C. Freiling

Department of Computer Science, Friedrich-Alexander University of Erlangen-Nuremberg, Am Wolfsmantel 46, 91058 Erlangen-Tennenlohe, Germany

ARTICLE INFO

Article history:

Received 9 May 2011

Received in revised form

8 June 2011

Accepted 11 June 2011

Keywords:

Memory forensics

Memory acquisition

Memory analysis

Live forensics

Microsoft windows

ABSTRACT

Traditional, persistent data-oriented approaches in computer forensics face some limitations regarding a number of technological developments, e.g., rapidly increasing storage capabilities of hard drives, memory-resident malicious software applications, or the growing use of encryption routines, that make an in-time investigation more and more difficult. In order to cope with these issues, security professionals have started to examine alternative data sources and emphasize the value of volatile system information in RAM more recently. In this paper, we give an overview of the prevailing techniques and methods to collect and analyze a computer's memory. We describe the characteristics, benefits, and drawbacks of the individual solutions and outline opportunities for future research in this evolving field of IT security.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

With the widespread use of computer systems and network architectures, digital cyber crime has, unfortunately, aggravated as well. According to a recent publication by the [Internet Crime Complaint Center \(2010\)](#), a partnership between the National White Collar Crime Center (NW3C) and the Federal Bureau of Investigation (FBI), the number of complaints filed to the institution has almost gone up by the factor 20 within less than a decade. In 2009, more than 336,000 reports about different types of illicit activity such as online fraud, identity theft, and economic espionage were registered. The yearly monetary loss of complaints referred to law enforcement was estimated to be nearly \$560 million. As a survey by the [Computer Security Institute \(2009\)](#) shows, companies may lose up to several hundred thousand dollars in the course of an incident. In such cases, a forensic investigation of the affected machines may prove helpful for reconstructing the

actions that led to the security breach, finding relevant pieces of evidence, and possibly taking legal actions against the adversary.

Traditional approaches in computer forensics mostly described the acquisition and analysis of *persistent* system data. Respected procedures usually involved powering off the suspect machine, creating an exact bit-by-bit image of the corresponding hard disks and other storage media, and performing a post-mortem examination of the collected information ([U.S. Secret Service, 2006](#); [U.S. Department of Justice, 2008](#)). Obtaining a copy of physical Random Access Memory (RAM) was, on the other hand, frequently neglected by first responders ([Shipley and Reeve, 2006](#); [Hoglund, 2008](#)), even though guidelines stressed the necessity of securing digital evidence with regard to the *order of volatility*, i.e., from the volatile to the less volatile, as early as in 2002 ([Brezinski and Killalea, 2002](#); [Casey, 2004](#); [Farmer and Venema, 2005](#)). In the face of ever-growing hard drive storage capabilities ([Oswald,](#)

* Corresponding author.

E-mail addresses: stefan.voemel@informatik.uni-erlangen.de (S. Vömel), felix.freiling@informatik.uni-erlangen.de (F.C. Freiling).
1742-2876/\$ – see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.diin.2011.06.002

2010), and correspondingly, tremendous efforts to analyze media in time (Mrdovic et al., 2009; Walters and Petroni, 2007; Shipley and Reeve, 2006) as well as a rising number of memory-resident malicious software applications (Moore et al., 2003; Rapid7 LLC, 2004; Sparks and Butler, 2005; Bilby, 2006), the restoration of transient and system state-specific information has, however, also moved more gradually into the focus of current research, beginning with the *Digital Forensic Research Workgroup* (DFRWS) challenge in 2005 (DFRWS, 2005).

This shift in practices has been driven and inspired by several other developments, too: First, “pulling the plug” on a company server may negatively affect productivity in certain cases and cause substantial losses due to unexpected down times. Furthermore, depending on the configuration of the system, file system journals may be damaged during the shutdown process, or the machine may be difficult to restart. Thus, there is a demand to minimize interferences with existing business and enterprise processes. Second, some programs are explicitly designed to make no or preferably as little persistent changes as possible on the hard disk of the user. Contemporary examples for this type of software are the *Mozilla Firefox* web browser with its private browsing capability (Mozilla Foundation, 2008; Aggarwal et al., 2010) or utilities included in the *PortableApps.com* project (Rare Ideas, 2010). For this reason, forensic analysts must adapt their strategies and also search in volatile system storages for traces and data remnants, including usernames, passwords, and text fragments. Moreover, many modern operating systems include support for file or even full disk encryption (Microsoft Corporation, 2009; Apple Inc., 2010; Saout, 2006). Similar functions are provided by freely-available open source tools, e.g., *TrueCrypt* (TrueCrypt Foundation, 2010), or commercial products such as *SafeGuard Easy* (Sophos Plc, 2010) or *PGP* (PGP Corporation, 2010). Because of the transparent design and ease of use of these software products, security professionals are likely to face an increasing number of encrypted drives that make traditional investigations infeasible (Getgen, 2009). Restoring a cryptographic key from memory might be the only possibility to get access to the protected data area in this case. The same holds true for packed malicious binaries. Malware writers typically employ compression, armoring, and obfuscation techniques to make reverse engineering and static analysis of their code more difficult (Sharif et al., 2009; Rolles, 2009; Brosch and Morgenstern, 2006; Young and Yung, 2004). Memory inspection is a viable solution to cope with these issues and extract the unpacked and decrypted executable directly from RAM.

As can be seen, a myriad of valuable information is stored in volatile memory that is usually lost when the target computer is powered off. Failing to preserve its contents may thus destroy a significant amount of evidence.

1.1. Motivation for this paper

Over the last 5 1/2 years, considerable research has been conducted in the field of memory forensics, and various methods have been published for capturing and examining the volatile storage of a target machine. However, many techniques solely apply to specific versions of operating

systems and architectures or only work under certain conditions. Moreover, depending on the technology used, the reliability and trustworthiness of generated results may vary. For these reasons, security professionals must have a thorough understanding of the capabilities and limitations of the respective solutions in order to successfully retrieve pieces of evidence and complete a case. A complete description of the current state of the art appears to be missing at the time of this writing though, restricting (research) activities in this area to a number of renowned experts.

In this paper, we give a comprehensive and structural overview of proven approaches for obtaining and inspecting a computer’s memory. We explain the technical foundation of existing tools and methodologies and outline their individual strengths and weaknesses. Based on these illustrations, security analysts and first responders may choose an adequate acquisition and analysis strategy. In addition, we give an extensive summary of the relevant literature. This review serves as a good starting point for own future studies.

Please note that our explanations refer to the product family of *Microsoft Windows* operating systems. We assume that due to their high popularity and dominant market position (Net Applications, 2010), investigators are particularly likely to face Windows-based machines in practice. In addition, as we will see, a deep knowledge of internal system structures is required to collect digital evidence from a volatile storage. Covering other platforms such as Linux or Mac OS is therefore out of the scope of this paper. Interested readers are referred to Movall et al. (2005) and Suiche (2010) for more information on these topics.

1.2. Outline of the paper

This paper is outlined as follows: In Section 2, we briefly describe the memory management process and give an overview of the most important data structures that are required for this task. Current techniques and methods for creating a memory image from the target system are presented in Section 3, followed by a detailed illustration of the different investigative procedures in Section 4. A special framework for memory analysis activities, *Volatility*, is subject of Section 5. We conclude with a summary of our work and indicate opportunities for future research in this area in Section 6.

2. Technical background

Modern multi-tasking operating systems typically do not access physical memory directly, but rather operate on an abstraction called *virtual memory*. This abstraction of physical RAM needs specific hardware support (the so-called *Memory Manager* or *Memory Management Unit*) and offers several inherent advantages, e.g., the possibility of providing each process with its own protected view on system memory as well as monitoring and restricting read and write activities with the help of privilege rules (Intel Corporation, 2011). The layout between the virtual and physical address space may differ though, and blocks of virtual memory do not necessarily map to contiguous physical addresses as illustrated in Fig. 1.

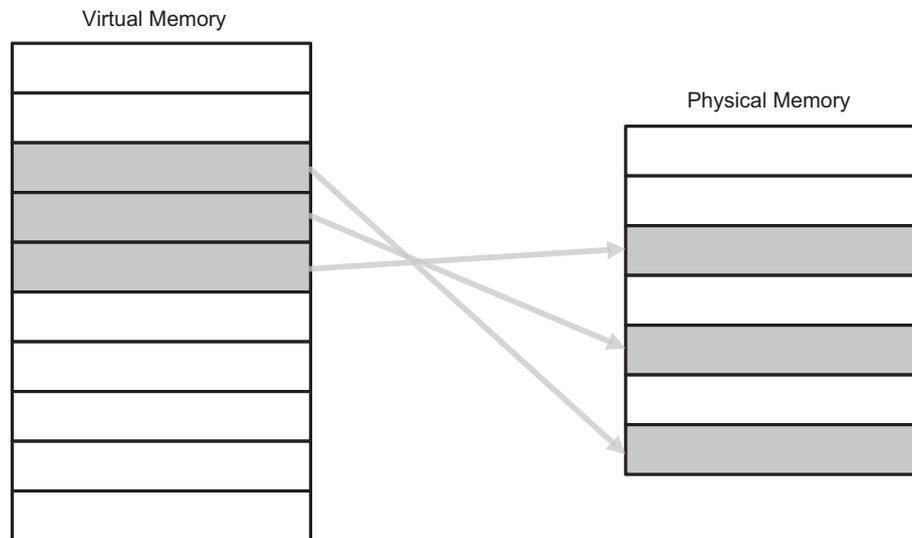


Fig. 1 – Mapping Virtual-to-Physical Memory (Source: Based on Russinovich et al., 2009, p. 15).

2.1. Memory address space layout

On Microsoft Windows operating systems, each process has its own *private virtual address space* that is non-accessible to other running applications unless portions of memory are explicitly shared. Thereby, collisions and privilege violations between different executables are prevented.

A 32-bit $\times 86$ user process is equipped with 2 GB of virtual memory by default (Russinovich et al., 2009). The other half of the address space (0×80000000 to $FFFFFFFF$) is reserved for system usage.¹ Kernel space is shared between and available to all system components. Thus, as Russinovich et al. (2009, p. 17) argue, it is vital that kernel-mode applications “be carefully designed and tested to ensure that they don’t violate system security and cause system instability”. Critical memory regions include, for example, the system pools that store volatile data that must not/may be paged out to hard disk as well as the process page tables that are required for virtual-to-physical address translation as we will see in the following section (see also Fig. 2).

2.2. Virtual address translation

As we have already explained, programs usually operate on virtual memory regions only. Therefore, to manipulate the respective physical data, the Memory Manager must continuously translate (*map*) virtual into physical addresses. This procedure works as follows (Intel Corporation, 2011; Russinovich et al., 2009): At the hardware level, volatile storage is organized into units called *pages*. A common size of such pages is 4 kB on $\times 86$ -platforms. To reference a page, the

¹ For reasons of simplicity, we assume that advanced memory management features such as large address spaces, Address Windowing Extension (AWE), and Physical Address Extension (PAE) are turned off. For more information on these concepts as well as details about the 64-bit address space layout, please refer to Russinovich et al. (2009).

operating system implements a two-level approach: For every process, the operating system maintains a *page directory* that saves pointers (*Page Directory Entries*, PDEs, 4 bytes each, containing a pointer and several flags) to 1,024 different *page tables*. Page tables, in turn, contain up to 1,024 links (*Page Table Entries*, PTEs, 4 bytes each) to the corresponding page in main memory. Thus, in order to translate a virtual to a physical address, the Memory Manager first needs to recover the base address of the page directory. It is stored in the CR3 register of the processor and is reloaded from the `_KPROCESS` block of the process at every context switch.² The first 10 bits of the virtual address can then be used as an index into the page directory to retrieve the desired PDE. With the help of the PDE and the *page table index*, i.e., the subsequent 10 bits of the virtual address, the page table and PTE in question are identified in the next step. To find the appropriate page and data in RAM eventually, the PTE and the 12-bit *byte index* of the virtual address are parsed. A summary of the entire process is presented in Fig. 3.

2.3. Paging

With respect to the address translation sequence outlined in the previous section, we have implicitly assumed that the requested information is always available in main memory. However, in some cases, the total amount of virtual memory that is consumed by running processes is larger than the size of the entire physical storage. To cope with these scenarios, the operating system needs to temporarily swap out (*page*) memory contents to hard disk in order to free required space. When a thread attempts to operate on a swapped-out page,

² The kernel process (`_KPROCESS`) block is part of a larger structure, the executive process (`_EPROCESS`) block, that serves as an internal representation for a Windows process. We will illustrate the `_EPROCESS` block in more detail in the memory analysis section of this paper. Further information can also be found in the work of Russinovich et al. (2009).

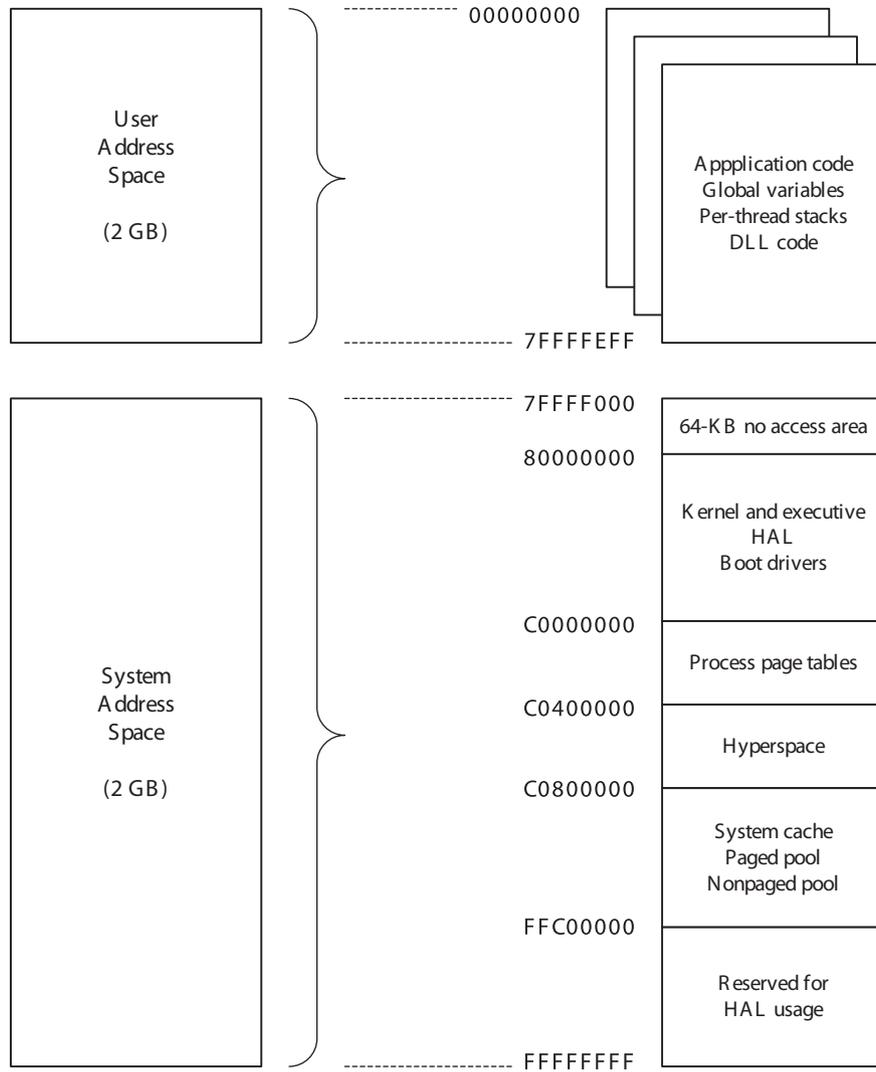


Fig. 2 – Virtual Address Space Layout (Source: Based on Russinovich et al., 2009, p. 738).

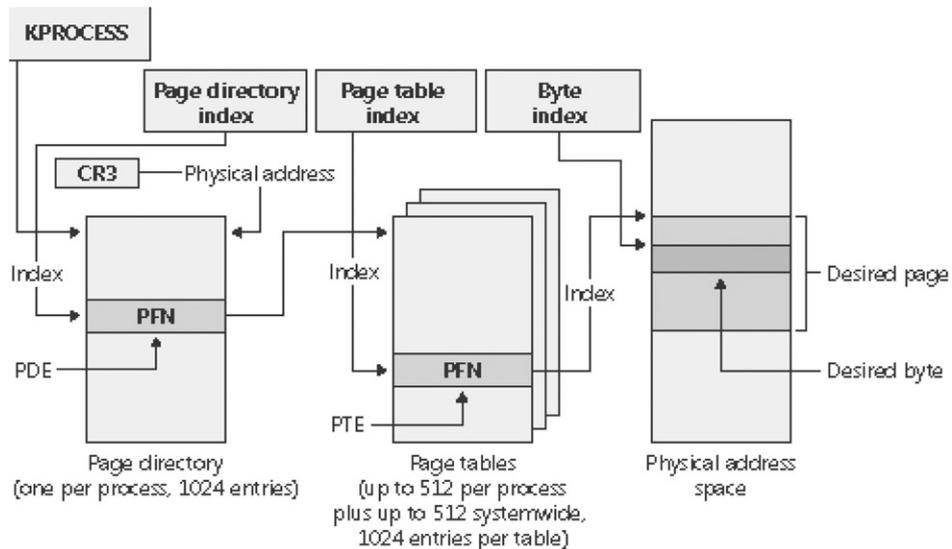


Fig. 3 – Virtual-to-Physical Address Translation (Source: Russinovich et al., 2009, p. 763).

the memory management unit generates a *page fault* interrupt that is handled by the operating system such that the requested information can be transferred back into RAM.

Whether or not a virtual address has been paged to disk is indicated by the *Valid* flag, the least significant bit, of a PDE or PTE. When it is set to 1, the entry is regarded as valid, and the respective data is accessible in memory. In contrast, when the flag is cleared and both the *Transition* (11) and *Prototype* (10) bit are set to 0, an address in a *page file* is referenced (see Fig. 4). The default page file `pagefile.sys` is saved in the root directory on the partition the operating system is installed on. However, in total, up to 16 different page files with a maximum size of 4,095 MB each can be supported on x86-based platforms. The name and location of the files are specified in the registry string `HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management\PagingFiles`, a field in the PTE of the virtual address, the *Page File Number* (PFN, bits 1 to 4), denotes the current file in use.

Please note that we have only given a brief overview of the flags and elements of a PDE and PTE, respectively. A more detailed explanation of the individual components and attributes can be found in the work of other authors (Russinovich et al., 2009; Savoldi and Gubian, 2008; Kornblum, 2007b; Maclean, 2006). The concepts described in this section need to be thoroughly understood though, because they form the basis for many memory acquisition and analysis tools presented in later parts of this paper.

3. Acquisition of volatile memory

Techniques for capturing volatile data are conventionally divided into *hardware-* and *software-*based solutions in the literature (e.g., see Vidas, 2006, 2010; Maclean, 2006; Garcia, 2007; Libster and Kornblum, 2008). While the latter ones depend on functions provided by the operating system, hardware-based approaches directly access a computer's memory for creating a forensic image and, therefore, have long been regarded to be more secure and reliable. A publication by Rutkowska (2007) indicates, however, that these assumptions no longer hold true. Moreover, several concepts that have been proposed more recently rely on a combination of both hardware and software mechanisms and cannot be clearly categorized with the existing terminology (Libster and Kornblum, 2008; Halderman et al., 2009). For this reason, we believe that a classification solely on implementation-specific attributes is obsolete and is not capable of properly characterizing the latest developments any more.

A more viable suggestion is assessing the different technologies with respect to the requirements that are necessary to obtain a (sound) memory copy of the target machine. Three major criteria identified by Schatz (2007a,b) are the fidelity of a generated image, the reliability of produced results, and the availability of the used mechanism. Fidelity enforces that the created image is “a precise copy [of] the original host's memory” (Schatz, 2007a, p. S128). It depends on the “atomicity” of the snapshot (which is not further defined) as well as the integrity of the host OS's memory and unallocated memory regions. Reliability dictates that either a trustworthy

result is produced or none at all. Finally, availability stipulates that the method must be “working on arbitrary computers (or devices)” (Schatz, 2007a, p. S128).

Inspired by this work, we define and adapt the following two factors *atomicity* and *availability* for our evaluation. The latter refers to the applicability of a certain technique on arbitrary system platforms for any given scenario. More precisely, a technique that is characterized by a high availability does not make any assumptions about particular pre-incident preparatory measures or pre-configurations and can be applied even without detailed knowledge about the execution environment. Atomicity, on the other hand, intuitively reflects the demand to produce an accurate and consistent image of a host's volatile storage. In more detail, an atomic snapshot is a snapshot obtained within an “uninterrupted” *atomic action* in the sense of a *critical section* as it is used in operating systems and concurrent programming (Lynch et al., 1993). It is free of the signs of concurrent system activity.

Based on these definitions we are able to broadly classify and compare the most common memory acquisition approaches. The corresponding decision matrix is visualized in Fig. 5.

Please note that the exact positioning of the methods within the fields of the matrix may certainly be subject to discussion at parts and is not the primary intention of the authors. Likewise, we acknowledge that the selected dimensions do not yet describe the characteristics of the different technologies in their entirety. In particular, our illustration does not adequately reflect the aspect of what Schatz' calls *reliability*, but what we would refer to as *correctness*, in compliance with forensic guidelines (e.g., see Scientific Working Group on Digital Evidence (SWGDE), 2009). A given procedure would therefore be regarded as “correct” if the content of the image file results from an assignment of the content of the target system's memory.³ Proving the correctness of this operation is a non-trivial task though, because results can generally not be reproduced due to the volatile nature of the collected information (see Sutherland et al., 2008).

In spite of the issues outlined above, we believe that our model serves as an initial attempt to formalize the memory acquisition process and is capable of giving investigators a good insight into when (not) to choose a specific solution. Vital points to keep in mind are especially:

- An ideal acquisition method is characterized by both a high atomicity and availability and is therefore located in the right upper corner of the matrix.
- Techniques that are listed in the right half of the matrix must generally be favored upon techniques that are grouped on the left side, because they are superior concerning atomicity (and possibly availability as well).

³ More formally, correctness means that the value of every individual bit of the created image file results from an assignment of the value of the corresponding bit of the target system's memory (and no other value). Note that since this definition does not refer to the specific time of the assignment, correctness does not imply atomicity (and vice versa).

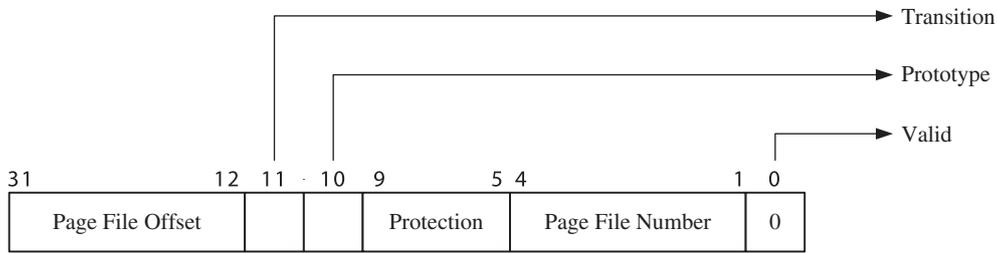


Fig. 4 – Invalid Page Table Entry (PTE) (Source: Based on Russinovich et al., 2009, p. 775).

- Methods that are located in the bottom field on the left side of the matrix (currently: none) are generally not suitable for obtaining volatile data in a forensically sound manner and must not be considered further.
- Approaches that are categorized in the right bottom field of the matrix are applicable for scenarios where an investigator has sufficient time for pre-incident preparation.
- Techniques that are listed in the right upper field of the matrix are especially suited for *smoking gun* situations, i.e., where little time between the incident and investigation phase has passed.

In the following, we describe the different approaches in detail and outline their individual benefits, peculiarities, and drawbacks.

3.1. Memory acquisition using a dedicated hardware card

One of the first propositions for obtaining a forensic image of a computer’s RAM storage was the use of a special hardware card. Carrier and Grand (2004) presented a proof-of-concept (PoC) solution called “Tribble” that makes use of Direct

Memory Access (DMA) operations to create a copy of physical memory, thereby bypassing possibly subverted structures of the operating system. The card is installed as a dedicated PCI device and is capable of saving volatile information to an attached storage medium. Upon pressing an external switch, the card is activated, and the imaging procedure is initiated. During this process, the CPU of the target host is temporarily suspended to prevent an attacker from executing malicious code and illegitimately modifying the status of the system. Once all operations are completed, control is given back to the operating system, and the acquisition card returns to an idle state again. Two less known but comparable implementations are proposed by Petroni et al. (2004) with their “Copilot” prototype and BBN Technologies (2006) in the form of “FRED”, the Forensic RAM Extraction Device.

As the described solutions do not rely on functions provided by the operating system, they are assumed to be generally suitable for acquiring an accurate image of volatile memory. Since the processor of the host in question is successfully halted, the imaging operation can be atomically completed without interference by other processes as well. Last but not least, because all information is directly retrieved from physical RAM, the procedure is believed to act outside

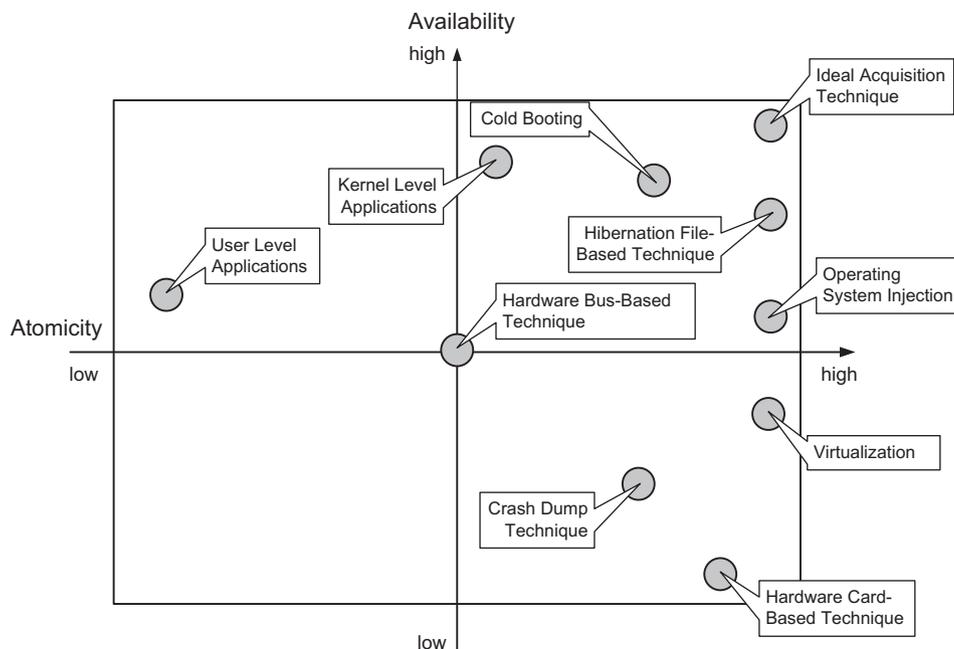


Fig. 5 – Classification of Memory Acquisition Techniques (Source: Based on Schatz, 2007b).

the view of malicious applications such as rootkits, resulting in a “true” picture of a target’s memory (Kornblum, 2006). Rutkowska (2007) has, however, recently proven that it is possible to present a different view of physical memory to the peripherals by reprogramming the chipset of the NorthBridge. Due to this innovation, several authors conclude that hardware cards can no longer be fully trusted and must not be regarded as forensically sound any more, making the development of more robust and reliable memory collecting techniques necessary (Libster and Kornblum, 2008; Ruff, 2008; Rutkowska, 2007).

In addition to these concerns, it is also important to emphasize that a PCI card must be set up *prior* to its use. This characteristic limits its application to special scenarios only. Carrier and Grand (2004, p. 12) clarify that “the device has not been designed for an incident response team member to carry in his toolkit”, but “rather needs to be considered as part of a forensic readiness plan”. According to the authors, a card is thus most beneficial when installed on (business-)critical servers, “where an attack is likely and a high-stake intrusion investigation might occur”. Alternative options for deployment are, for instance, within a honeypot environment.⁴ Thereby, it is possible to capture volatile information after a machine has been compromised and learn more about the tools, tactics, and motives of adversaries.

3.2. Memory acquisition via a special hardware bus

As an alternative to PCI cards, several authors suggest reading volatile memory via the IEEE 1394 (FireWire) bus (Dornseif and Becher, 2004; Becher et al., 2005). A corresponding forensic-related application is described by Piegdon and Pimenidis (2007). While the original code targets Linux and Mac OS platforms, Boileau (2006b, 2008), Panholzer (2008), and Böck (2009) demonstrate the feasibility of the approach for different versions of Microsoft Windows. Ruff (2008, p. 84) adds that “any hardware bus can potentially be used for physical memory access”. For instance, a proof-of-concept utility that illustrates Direct Memory Access (DMA) operations with the help of the PCMCIA (PC Card) bus was published at the ShmooCon 2006 conference (Hulton, 2006). The sample code has, however, not been published yet.

Retrieving volatile information via the IEEE 1394 bus can address some of the issues we have outlined in the previous section. For instance, the interface is present by default in a great part of systems, especially laptops. However, as Vidstrom (2006) points out, the use of the technique may cause random system crashes and similar reliability problems when accessing regions in the Upper Memory Area (UMA). Other authors have also indicated inconsistencies after comparing created images with raw memory dumps (Carvey, 2007; Boileau, 2006a). For this reason, similar to the hardware card-oriented approach illustrated above, FireWire-based techniques are not sufficiently reliable to obtain a precise copy of a computer’s RAM.

⁴ A honeypot is “an information system resource whose value lies in unauthorized or illicit use of this resource” (Spitzner, 2003b,a). It acts as an electronic decoy for studying the behavior of (Internet) miscreants (HoneyNet Project, 2004).

3.3. Memory acquisition with the help of virtualization

With the help of virtualization, it is possible to emulate complete, isolated, and reliable system environments, so-called *virtual machines*, on top of a host computer (Smith and Nair, 2005). A special software layer, the *virtual machine monitor* (VMM), is responsible for sharing as well as managing and restricting access to the available hardware resources. By emulating replicas of the different physical components, each virtual machine is equipped with its own virtual processor, memory, graphics adapter, network and I/O interface and may run in parallel to other *guest* systems.

One exceptional characteristic of a virtual machine is its capability to be suspended, i.e., to pause its execution process. Thereby, the state of the guest operating system is temporarily frozen, and its virtual memory is saved to hard disk on the underlying host. For instance, in the case of a VMware-based machine, all volatile data is saved to a `.vmem` file located in the working directory of the virtual machine (VMware, Inc. 2010). Thus, by simply copying the respective file, an atomically-generated snapshot of main memory can be easily retrieved.

In 2007, Carvey (2007, p. 95) pointed out that “virtualization technologies do not seem to be widely used in systems that require the attention of a first responder”. With the growing importance of Internet-hosted services, this situation is likely to change though, and investigators will increasingly have to examine incidents on virtual machines. With respect to these scenarios, the memory acquisition approach is forensically sound and only requires little effort.

3.4. Memory acquisition using software crash dumps

All versions of Microsoft Windows 2000 and above can be configured to write debugging information, so-called *memory dump files*, to hard disk when the machine unexpectedly stops working (Microsoft Corporation, 2010e). In case of a critical failure, the system state is frozen, and the main memory as well as relevant CPU information are saved in the system root directory for later examination. Preserving the contents of processor registers during this operation is a unique characteristic of this method (Russinovich et al., 2009; Carvey, 2007).⁵ A dump file can then either be opened with the *Microsoft Debugging Tools for Windows* (Microsoft Corporation, 2010a) or be manually analyzed as explained by Schuster (2006a, 2008a).

With regard to a forensic investigation, the system service may be intentionally interrupted by using a third-party application (Microsoft Corporation, 2010g; Open System Resources, 2009) or the built-in *CrashOnCtrlScroll* feature. In the latter case, the dump is generated upon pressing the keyboard shortcut `Right Ctrl + ScrollLock + ScrollLock` (Microsoft Corporation, 2010d). It is important to note that this capability is disabled by the operating system by default though. Activation requires editing a certain value in the Windows registry (Russinovich et al., 2009), a subsequent

⁵ We assume the system is configured to save at least a full kernel memory dump. For more information on this subject, please refer to Microsoft Corporation (2010e).

reboot,⁶ and only works with PS/2 keyboards initially. For universal serial bus (USB) devices, a separate hotfix must be installed (Microsoft Corporation, 2010f). Due to these modifications, the applicability of the technique is forensically only suitable in specific situations, analogously to the use of a dedicated PCI card as explained in Section 3.1. Moreover, even though the memory snapshot is atomically created by halting the processing unit, parts of the system page file are overwritten during this operation (Russinovich et al., 2009). Compared to other approaches, the acquisition technique is therefore more invasive and leads to a higher degree of contamination on the target system.

3.5. Memory acquisition with user level applications

Especially at the beginning of memory-related forensic research, various software solutions have been published by different third parties to acquire a copy of physical memory directly from user space. A prominent example is *Data-Dumper* (dd) by Garner (2009) which is part of the *Forensic Acquisition Utilities* (FAU) suite. When run with administrative privileges, dd invokes the internal `\\.\Device\PhysicalMemory` section object to create a full memory dump of the target machine (Crazylord, 2002). Due to security reasons, user level access to this object was restricted in Windows Server 2003 and later operating systems though (Microsoft Corporation, 2010b). With regard to the configuration of modern computer platforms, the technique is thus only of very limited value to date.

Rather than creating a (much larger and more time-consuming) forensic duplicate of the entire volatile storage of a host, several other utilities permit saving the address space of a single process: With *PMDump*, it is possible to “dump the memory contents of a process to a file” (Vidstrom, 2002). The *Process Dumper* utility (PD) by Klein (2006b) obtains a process’s environment and state by retrieving the data and code mappings of an executable. Since all collected information is redirected to standard output by default, the final image can be easily transferred over a remote connection for further investigation, e.g., with *netcat* (nc), a freely-available network administration utility (Craton, 2009).⁷

Both *PMDump* and *Process Dumper* also have various drawbacks though: First, they are closed source and use a proprietary data format that make the development of additional parsing tools more difficult. Second, as the programs require the specification of a process ID, a corresponding process listing utility must be run in the first place. This operation further affects the level of contamination on the target host, however.

As the techniques described above neither rely on specific hardware setups or devices nor do they require a pre-defined system configuration, they are generally suitable for most incident scenarios and permit capturing a forensic image even

in case a first responder only has little time for preparation. However, while the imaging process is comparatively easy to perform, the approaches still suffer from various inherent weaknesses: First, as we have already indicated, many programs only work on specific operating systems or are bound to specific architectures, e.g. the x86 32-bit platform. Second, as all applications must be loaded into memory before they can be executed, valuable information may be destroyed before it can be preserved. Sutherland et al. (2008) have shown in a study that the impact of the tools on a target machine can be significant. What is worse, other processes are usually not impeded from altering the volatile storage during the imaging operation, resulting in a potentially non-atomic and “fuzzy snapshot” (Libster and Kornblum, 2008, p. 14) of the source data. Last but not least, as all methods depend on functions provided by the operating system, they are vulnerable to subversion. A rootkit might, for instance, deny direct access to the physical memory object or return a modified representation of RAM during the image generation operation in order to evade detection (Vidas, 2006; Bilby, 2006; Sparks and Butler, 2005). Although these types of manipulations are likely to indicate the presence of a malicious software application to the eyes of a skilled analyst (Kornblum, 2006), relying on an untrusted operating system eventually “decreases the reliability of the evidence” (Carrier and Grand, 2004, p. 6), too.

3.6. Memory acquisition with kernel level applications

To mitigate the limitations of user space acquisition utilities described in the previous section, software vendors increasingly provide kernel level drivers to create a forensic image of physical memory. Freely-available distributions include Mantech’s *Memory DD* (mdd, ManTech CSI, Inc., 2009), MoonSols’ *Windows Memory Toolkit* (MoonSols, 2010), and *Memoryze* by Mandiant (Mandiant, 2010). Commercial alternatives are offered by Guidance Software (*WinEn* which is a part of *EnCase*, Guidance Software (2010)), *AccessData* (Forensic Toolkit, AccessData (2010)), *GMG Systems* (*KnTDD* which is part of the *KnTTools*, GMG Systems, Inc. (2007)), and *HBGary* (*Fastdump Pro*, HBGary (2009)). Even though the different solutions are usually available for various versions of the Windows product family, they still cannot overcome the inherent weaknesses we have already outlined. These include manipulations due to concurrently running (system) processes and, consequently, a not sufficiently precise representation of a host’s memory as well as the possibility of falling prey to a compromise attempt (see Section 3.5).

In order to cope with these concerns, Libster and Kornblum (2008) propose integrating the capturing mechanism as an operating system module into the system core which is loaded at boot time and invoked by a special keyboard trigger. According to the authors, distinctive characteristics of the module would be the capability to halt active system processes, thereby ensuring atomic operations. In addition, support for several storage dump locations, e.g., a remote network drive or an externally attached media, is encouraged. For increased security, the use of hardware-side, read-only memory flags or encrypted Trusted Platform Modules (TPMs (Group, 2007)) is suggested. These guarantee the integrity of the imaging operation and prevent attacks on the executing code.

⁶ A technique for omitting the mandatory reboot of the operating system was outlined by Ruff (2008).

⁷ Netcat does not establish an encrypted communication channel by default. To securely transfer data over the network, the connection can be tunneled over the SSH (*Secure Shell*) protocol. Alternatively, the *cryptcat* implementation may be used as well. For more information on these topics, please see Farmer and Venema (2005).

3.7. Memory acquisition via operating system injection

Schatz (2007a) has introduced a proof-of-concept application called *Body-Snatcher* which injects an independent operating system into the possibly subverted kernel of a target machine. By freezing the state of the host computer and solely relying on functions provided by the acquisition OS, an atomic and reliable snapshot of the volatile data may be created. The presented prototype is, however, platform-specific to a high degree due to its low level approach and high complexity. In addition, it is limited to single processor mode at the time of this writing, consumes a significant amount of memory, and only supports the serial port for I/O operations. Therefore, even though the concept is very promising, its technical constraints still need to be resolved, before it can be truly applied in real-world situations.

3.8. Memory acquisition via cold booting

Another encouraging approach has been outlined by Halderman et al. (2009). It is based on the observation that volatile information is not immediately erased after powering off a machine, but may still be recovered for a non-negligible amount of time (Chow et al., 2005). By artificially cooling down the respective RAM modules, e.g., with liquid nitrogen, remanence times may even be substantially prolonged. The target machine can then be restarted (*cold booted*) with a custom kernel to access the retained memory in the next step. The usability of this approach has been proven in a number of recent works: Vidas (2010) has published *AfterLife*, a proof-of-concept demonstration that copies the contents of physical RAM to an external storage medium after rebooting. Chan et al. (2008, 2009) implemented a special booting device that revives the host system environment after cutting power and provides the investigator with an interactive shell to retrieve state-related system data. Most of the projects are still in a state of flux though and are not designed to be used in daily practice yet.

3.9. Memory acquisition using the hibernation file

The Windows hibernation file (`hiberfil.sys`) is increasingly regarded as another source of valuable information (Carrier and Grand, 2004; Schatz, 2007a; Libster and Kornblum, 2008; Ruff, 2008; Zhao and Cao, 2009). It is stored in the root directory on the same partition Microsoft Windows is installed on and contains runtime-specific information. When the operating system is about to enter an energy-saving sleep mode and is suspended to disk, the system state is frozen, and a snapshot of the existing working set is preserved. Thus, in contrast to many other software-based solutions we have described, investigators are able to capture a consistent and atomically-generated image of volatile data. However, it is important to note that the file cannot be used to reproduce the state of physical memory in its entirety as “for hibernation, the goal is to intelligently decide which pages are saved in the hibernation file” (Russinovich et al., 2009, p. 840). Hence, the quantity and quality of extracted pieces of evidence may be limited. In addition, the hibernation file is compressed in order to save disk space and uses a proprietary format which

has long impeded thorough understanding and applicability in forensic investigations. Details about its structure have been discussed in more recent publications though (Suiche, 2008a). A working prototype that is capable of converting the file to raw dump format as well as reading and writing specific memory sections was developed in the course of the *SandMan* project (Ruff and Suiche, 2007; Suiche, 2008b). At the time of this writing, it has been superseded by *MoonSols* (MoonSols, 2010). The commercial version supports the entire family of Windows operating systems, including both 32- and 64-bit versions.

4. Analysis of the acquired memory image

After a forensic copy of physical memory has been generated with one of the techniques outlined in the previous sections, an in-depth analysis of the acquired data can begin. Primitive approaches that are described in the literature rely on simple string searches, e.g., with command line utilities such as `strings` and `grep` or more powerful applications such as *WinHex* (X-Ways Software Technology AG, 2010), to look for suspicious patterns, usernames, passwords, and other textual representations in the created image (Stover and Dickerson, 2005; Zhao and Cao, 2009). While these methods are easy to apply, they are also noisy, cause a huge overhead, and lead to a large number of false positives (Beebe and Dietrich, 2007). Beebe and Clark (2007, p. S49) argue that “[f]requently, investigators are left to wade through hundreds of thousands of search hits for even reasonably small queries [...] – most of which [...] are irrelevant to investigative objectives”. Even when string searches produce quite accurate results, they typically do not take into account the context of the respective information and, as a consequence, are of limited help to the investigation process. For instance, as Savoldi and Gubian (2008, p. 16) point out, “the retrieval of a potentially compromising string (e.g. ‘child porn’) certainly provides evidence if found in the memory assigned to a process l[a]unched by the user, but it would be likely rejected by jury if that memory belonged to a piece of malware”. In addition, Hejazi et al. (2009, p. S123) note that the “existence of unknown sensitive data in the memory is the main important limitation of this method”. Thus, a forensic analyst may, for example, solely look for specific keywords, but at the same time, disregard “names, addresses, user IDs, and strings that are not present in the list the investigator is looking for while they are present in the memory dump and are of paramount importance for the investigative case” (Hejazi et al., 2009, p. S123). In sum, these procedures are not sufficiently reliable with regard to a forensic investigation. Efficient solutions to these problems have been developed over the last years but are not further illustrated in the remainder of this paper since they fall in the area of information retrieval (IR) and text mining tasks. Good introductions to these topics have been compiled by Beebe and Clark (2007) as well as Roussev and Richard III (2004) though.

As a viable alternative to string searching algorithms, security professionals recommend a more structured methodology to find valuable traces in memory. It involves examining *what types* of data may be contained within a captured



Fig. 6 – Subversion of a Process List Using Direct Kernel Object Manipulation (DKOM). (a) Structure of the Process List Before a DKOM Attack. (b) Structure of the Process List After a DKOM Attack.

image, how these types are defined, and where they are located. Relevant pieces of information are generally stored in the system or user address space, either directly in RAM or in the local page file, and include (Hoglund, 2008; Sutherland et al., 2008)

- the list of running system processes and possibly installed malicious software applications,
- cryptographic keys,
- the system registry,
- established network connections and network-related data such as IP addresses and port information,
- open files, and
- system state- and application-related data such as the command history, date and timestamp information.

The prevailing techniques to recover the specified artifacts are subject of the following sections.

4.1. Process analysis

An essential part of a forensic investigation is checking the integrity of the operating system on the suspect machine and distinguishing legitimate components from suspicious and potentially malicious applications. Thus, the identification

and sound examination of running processes becomes “a basic security task that is a prerequisite for many other types of analysis” (Dolan-Gavitt et al., 2009, p. 2).

Early approaches attempted to parse internal system resources to enumerate the list of loaded programs (Burdach, 2005; Betz, 2006; Garner, 2007). However, as various authors point out, an increasing number of malware executables use so-called rootkit techniques and subvert integral system structures to avoid detection (Aquilina et al., 2008; Ligh et al., 2010). For instance, the *FU* rootkit (Butler, 2005) implements a method known as *Direct Kernel Object Manipulation* (DKOM) to unlink itself from the `ActiveProcessLinks` list, a member of the `_EPROCESS` (executive process) block every process is represented by (Bilby, 2006). This operation is illustrated in Fig. 6(b), the skull icon over the second structure indicates the malicious rootkit process. Due to these interventions, list and table walking solutions are likely to fail and must not be seen as reliable.

In order to cope with these issues, Schuster (2006d) developed a signature-based scanner. It uses a set of rules that precisely describe the structure of a system process or thread, respectively. A sample signature is illustrated in Listing 1. In the given example, it is checked that the page directory is aligned at a page boundary, the control structures of the working thread are contained in kernel space, and the fields of the `DISPATCHER_HEADER` structure match pre-defined values.

The results of the scanner can be compared with the output of the standard process list in the next step. Differences and anomalies potentially indicate the presence of a malicious program and need to be inspected in more detail. Similar strategies are proposed by Walters and Petroni (2007), Carr (2006), and Rutkowska (2005). Walters and Petroni (2007, p. 14) note, however, that “a reliable pattern may rely on characteristics of an object that are not essential”. Consequently, an adversary may change the value of a *non-essential* field of the process object without affecting the stability of the underlying operating system. For instance, in the case of Listing 1, it is also possible to set the value of the `Size` field to zero. Thereby, the rule of the scanner is circumvented, and the respective process becomes invisible. To prevent such scenarios, Dolan-Gavitt et al. (2009) have created so-called *robust signatures* which are solely based on fields that are critical to system functionality. A manipulation of such a field leads to an automatic system crash and, thus, renders an attack useless. With regard to the `core_EPROCESS` structure outlined above, 72 fields (out of 221) suffice this requirement and, therefore, form strong candidates for a process signature.⁸

An alternative method is suggested by Zhang et al. (2009, 2010). The authors use a combination of scanning and list traversing techniques that rely on the *Kernel Processor Control Region* (KPCR). This region stores processor-specific data and contains a separate block, the *KPRCB*, that, for instance, saves CPU-related statistics as well as scheduling information about the current and next thread (Russovich et al., 2009). With the help of this information, the corresponding process list can be restored as follows (Barbosa, 2005; Ionescu, 2005; Dolan-Gavitt, 2008b): The KPCR includes a field `KdVersionBlock` that points to a `_DBGKD_GET_VERSION64` structure that, in turn, references the undocumented `_KDDEBUGGER_DATA64` structure. Part of this structure is the global variable `PSActiveProcessHead` that serves as a pointer to the list of active `_EPROCESS` structures.

In Microsoft Windows XP, both the KPCR and KPRCB are located at fixed addresses (`0xFFDF000` and `0xFFDF120`, respectively). In later versions, the base addresses are dynamically computed. Since the KPCR structure is self-referencing (offset `0x1C`) and the KPRCB starts at offset `0x120`, a simple signature can be generated though to discover the structures and enumerate the running programs in the next step (Aumaitre, 2009).

Listing 1 : Example of a process signature (Schuster, 2006d)

```
PageDirectoryTable != 0
PageDirectoryTable % 4096 == 0
(ThreadListHead . FLink > 0x7FFFFFFF) &&
(ThreadListHead . Blink > 0x7FFFFFFF)
DISPATCHER_HEADER . Type == 0x03
DISPATCHER_HEADER . Size == 0x1b
```

4.2. Cryptographic key recovery

Due to the spread of (freely) available encryption technologies, security professionals are likely to increasingly encounter

secured and password-protected files and disks in the future (see Section 1). In case an investigator is unable to obtain the desired information from the suspect, e.g., through social engineering, the restoration of cryptographic keys from volatile memory may therefore become a central task in the course of a forensic analysis.

To solve this problem, different methods are proposed in the literature: Hargreaves and Chivers (2008) describe a linear memory scanning technique that moves through RAM one byte at a time, using a block of bytes as the possible decryption key for the volume in question. The brute force-oriented procedure does not require a deep understanding of the underlying operating system and, thus, can be easily generalized according to the authors. However, it cannot be directly applied if the key is split, i.e., is not stored in a contiguous pattern in memory.

Shamir and van Someren (1999) seek for sections of high entropy to locate an RSA key within “gigabytes of data”. The solution exploits the mathematical properties of the cryptographic material. In contrast, the attack described by Klein (2006a) is based on the observation that both private keys and certificates are stored in standard formats. Klein (2006a) constructs a simple search pattern to easily dump the secret information from RAM. A pattern-like approach is also implemented by Kaplan (2007). His idea stems from the fact that, for reasons of security, cryptographic keys are generally not paged out to disk and, as a result, are typically saved in the *non-paged pool* of the operating system (see Section 2.1). The non-paged pool consists of a range of virtual addresses that always reside in physical memory (see also Russovich et al., 2009; Schuster, 2006b, 2008b). Since in kernel space, memory is a shared resource, allocated regions may be associated with an identifier, the so-called *pool tag* (Microsoft Corporation, 2010c). Consequently, “a cryptosystem-specific signature, consisting of the driver specific pool tag and pool allocation size are all that is necessary to extract pool allocations containing key material from a memory dump with an acceptably small number of false positives” (Kaplan, 2007, p. 18). As an alternative to the non-paged pool, it is possible to prevent regions of virtual address space from being written out to disk by calling the `VirtualLock` function. Thereby, “[p]ages that a process has locked remain in physical memory until the process unlocks them or terminates” (Microsoft Corporation, 2010h). Whether these mechanisms are relevant to the forensic recovery of cryptographic keys has not been investigated to the best knowledge of the authors though.

Walters and Petroni (2007) outline a different concept that relies on an analysis of publicly available source code. They identify internal data structures that are responsible for holding the master key. As Maartmann-Moe et al. (2009, p. S133) point out, Walter and Petroni “do, however, not describe how to locate the different structures in memory, and neither do they discuss the fact that some of these may be paged out, thereby breaking the chain of data structures that leads to the master key if only the memory dump is available for analysis”.

Last but not least, Halderman et al. (2009) suggest parsing a computer’s memory for key schedules. The authors leverage remanence effects in DRAM modules and launch a *cold boot* attack on the target machine (see Section 3.8). After loading a custom operating system, the volatile data is extracted, the

⁸ After performing a fuzzing test, the number of strong signature candidates was reduced from 72 fields to 43 (Dolan-Gavitt et al., 2009).

key material is retrieved, and the hard drive automatically decrypted. The performance of this process is improved in the works of Heninger and Shacham (2009). Likewise, Tsow (2009) presents an algorithm that is capable of recovering cryptographic information from memory images that are significantly decayed and is magnitudes faster than the original method. Finally, Maartmann-Moe et al. (2009) extend the conducted research on additional ciphers and illustrate the vulnerability of several well-known whole-disk and virtual-disk encryption utilities for these types of side-channel operations. However, it is important to note that these attacks can be mitigated by implementing cryptographic algorithms entirely on the microprocessor. A corresponding proof-of-concept application has been published by Müller et al. (2010) and has been further sophisticated more recently (Müller et al., 2011).

4.3. System registry analysis

The Windows registry is a central, hierarchically-organized repository for configuration options of the operating system and third-party applications (Microsoft Corporation, 2008). It is internally structured into a set of so-called hives, i.e., discrete, treelike databases that hold groups of registry keys and corresponding values (see also Russinovich et al., 2009). Most registry hives (e.g., HKLM\SYSTEM, HKLM\SOFTWARE) are persistently stored in the `system32\config` folder of the operating system.⁹ However, a number of volatile hives (e.g., HKLM\HARDWARE) are entirely maintained in RAM only and are created every time the system is booted.

While the examination of on-disk registry data is a quite established procedure for a forensic investigation to find possible pieces of evidence (Carvey, 2005; Mee et al., 2006; Chang et al., 2007), a solely memory-based approach has been documented by Dolan-Gavitt (2008c) only recently. In the following, we give a short overview about the prevailing techniques used in his work.

A registry hive consists of a so-called *base block* and a number of *hive bins* (hbins). The base block with a fixed size of 4 kB defines the start of the hive and contains a unique signature (`regf`) as well as additional meta information such as a timestamp that saves the time of the last access operation, an index to the first key, the `RootCell`, the internal file name of the hive, and a checksum. A hive bin is typically 4 kB wide (or a multiple of it) and serves as a container for *cells* that, in turn, store the actual registry data. Thus, to read a certain key or value from RAM, the correct hive and corresponding cell must be found first. The former task can be quite easily solved by creating a hive-specific signature: Internally, a hive is represented by a `_CMHIVE` structure which is allocated from the paged pool of the system (see Section 2.1) and referenced by the tag `CM10`. Furthermore, it embeds a sub-structure `_HHIVE` with the string constant `0xbee0bee0`. These pieces of information are sufficient to construct a unique search pattern and locate a hive in memory. To enumerate the entire list of loaded hives, the `HiveList` member structure must be

⁹ User-specific settings are saved in the file `NTuser.dat` that is located in the `%SystemDrive%\Documents and Settings\` folder.

simply followed in the last step. It acts as a link between the individual registry repositories.

Retrieving pre-defined keys or values from a memory image is slightly more complex: Dolan-Gavitt (2008a) notes that because “space for the hives is allocated out of the paged pool, there is no way of guaranteeing that the memory allocated for the hive will continue to be contiguous”. For this reason, a strategy is used that is similar to the virtual-to-physical address translation mechanism we have described in Section 2.2. A cell is generally referenced by a *cell index*. This field can be split into four different components. The first element, a one-bit flag, determines the stable (on-disk) or volatile *storage map* for the hive which are both saved in the `_HHIVE` sub-structure. With the help of the 10-bit directory index, the correct entry in the cell map directory can be located. The cell map directory refers to a cell map table with 512 (2⁹) entries that, in turn, point to the virtual address of the target bin. Using the third and fourth component of the cell index, the correct cell and cell information can be finally discovered (see Fig. 7). Dolan-Gavitt (2008e) has published a proof-of-concept utility that is capable of extracting the list of open keys and displaying the corresponding registry data. The software is freely-available for download and integrated into the Volatility analysis framework that we will introduce in more detail in a later section of this paper. It is, however, important to emphasize that the techniques described above only supplement traditional investigation methods. As later versions of Microsoft Windows map only 16-KB portions of a hive into RAM when they are needed (Russinovich et al., 2009), it is possible that “parts of the registry may have never been brought into memory in the first place” and, consequently, “it cannot be assumed that the data found in memory is complete” (Dolan-Gavitt, 2008c, p. S30).

4.4. Network analysis

Analyzing open network connections as well as inspecting incoming and outgoing network traffic is an integral part of a forensic investigation and becomes particularly important in the face of a potentially compromised system. Malicious applications typically bind to pre-defined ports and enable attackers to execute arbitrary commands on the target machine, disable security protections, and upload or download files (Aquilina et al., 2008; Ligh et al., 2010). Contemporary examples also include launching *Distributed Denial of Service* (DDoS) attacks and consuming the resources of a host in order to degrade its performance (Peng et al., 2007). A vast number of utilities, either for live response (e.g., TCPView (Russinovich, 2010), FPort (Foundstone, Inc., 2000)) or post-mortem examination (e.g., PyFlag (Cohen, 2008)), have been developed to reveal such threats. Since many network-related data are also temporarily stored in RAM, memory forensic techniques can help correlate results with the tools described above in order to gain a thorough understanding of an incident.

The first solution was suggested by Schuster (2006b). He scans the non-paged pool of the operating system to find allocations for listening sockets. His algorithm is based on a unique pool tag (`TCPA`, saved in little-endian format) and a pre-defined pool size (368 bytes) that can be recovered after disassembling the `tcpip.sys` driver where the respective

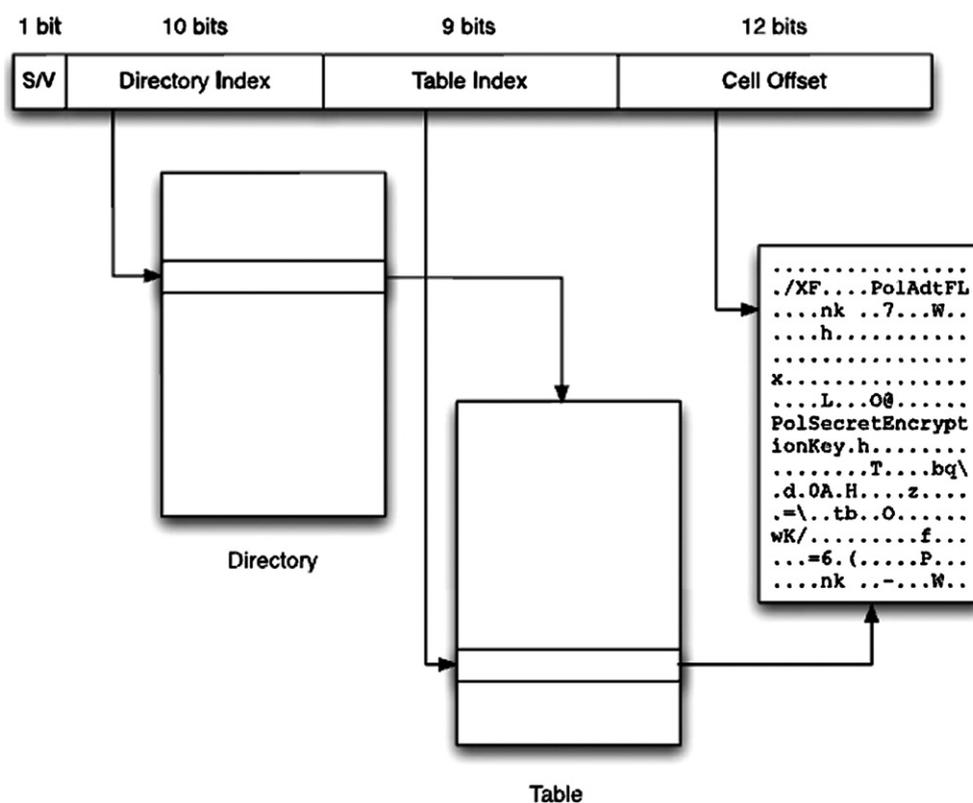


Fig. 7 – Structure of a Cell Index (Source: Dolan-Gavitt, 2008c, p. S28).

operating system functions are implemented in. Once the appropriate addresses have been identified, the socket list can be easily created. With a similar procedure, it is possible to retrieve the list of open network connections.

A different methodology is described by Ligh et al. (2010) and Okolica and Peterson (2010): The authors locate two internal hash tables (`_AddrObjTable`, `_TCBTable`) in the `tcpip.sys` driver file. Each of the hash tables references a singly linked list of objects that include information about the IP address and port bindings.¹⁰ By traversing the lists as shown in Fig. 8, existing sockets and open network connections can be enumerated.

It is interesting to note that a malicious manipulation of these lists, similar to the DKOM attack described in Section 4.1, appears to hamper system functionality and, thus, is likely to be detected in practice. For demonstration purposes, Ligh et al. (2010) manually unlinked an object to hide the presence of a listening socket. This operation disrupted the communication process though, and connections could no longer be initiated. From a forensic point of view, list crawling-based approaches can therefore be seen as reliable to date and permit a legit and unaltered view of network activities.

4.5. File analysis

Similar to the examination of running system processes (see Section 4.1), inspecting the list of open files and dynamically

¹⁰ Details about the structure of these objects can be found in the work of Ligh et al. (2010).

loaded libraries (DLLs) a program is referencing may be crucial to reveal suspicious activities in the course of an investigation. For instance, an adversary might inject a malicious DLL in the address space of a legitimate process by exploiting a remote vulnerability to hide her presence on the target machine (Miller and Turkulainen, 2006; Walters, 2006). As Walters (2006, p. 2) points out, these techniques “thwart conventional forms of file system forensic analysis” and “are commonly labeled as ‘anti-forensic’, since data is not written to any non-volatile storage and [...] would typically be lost during common incident response procedures”. To counter such threats, security professionals recommend analyzing the *Process Environment Block* (PEB), a structure which lives in the user address space and is part of every process (see Russinovich et al., 2009). The Process Environment Block contains a `Ldr` member with three doubly linked lists that save the full name, size, and base address of all loaded libraries.¹¹ Simply enumerating the individual lists may thus suffice to identify an injection attack. It is critical to stress though that the approach can be subject to manipulation as various rootkit authors have demonstrated (Darawk, 2005; Kdm, 2004). For this reason, Dolan-Gavitt (2007a) proposes an alternative methodology which is based on Virtual Address Descriptors (VADs).

¹¹ The three doubly linked lists (`InLoadOrderModuleList`, `InMemoryOrderModuleList`, `InInitializationOrderModuleList`) contain the same modules but in different order. For more information, please refer to Russinovich et al. (2009).

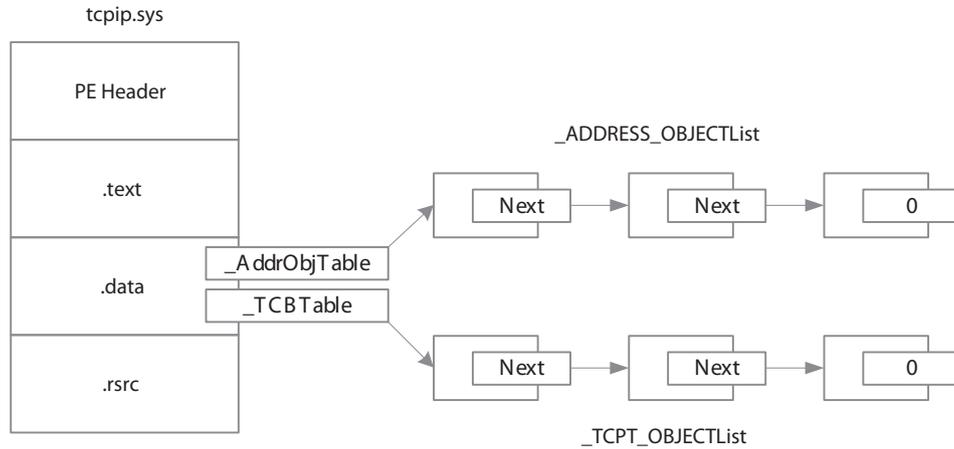


Fig. 8 – Enumerating the List of Socket and Connection Objects (Source: Based on Ligh et al., 2010, p. 681).

A *Virtual Address Descriptor* is a kernel data structure that is maintained by the memory manager to keep track of allocated memory ranges (Russovich et al., 2009). It stores information about the start and end of the respective addresses as well as additional access and descriptor flags. Each time a process reserves address space, a VAD is created and added to a self-balancing tree for maintenance reasons. A node in the tree is associated with a pool tag and is of type `_MMVAD_SHORT` (“VadS”), `_MMVAD` (“Vad”), or `_MMVAD_LONG` (“Vadl”) (Dolan-Gavitt, 2007a). The latter two store a pointer to a `_Control_Area` structure that, in turn, points to a `_File_Object` that holds the unique file name (see the lower half of Fig. 9). Consequently, the entire list of loaded modules can be retrieved by traversing the VAD tree from top to bottom and following the corresponding `_Control_Area` and `_File_Object` references.¹² Furthermore, it is possible to reconstruct the contents of memory-mapped files, most importantly, executables. While it is unlikely to completely rebuild the exact form of a binary due to changes in code variables at runtime (Kornblum, 2007a), the recovered copy can frequently be reverse engineered and inspected upon malicious behavior (see also Schuster, 2006c). These operations supplement traditional file carving techniques that are implemented in utilities such as *Foremost* (United States Air Force Office of Special Investigations, 2001) and *Scalpel* (Richard III, 2006) and are suited for scenarios where an on-disk restoration has failed. A proof-of-concept application has been published by Dolan-Gavitt (2007b) and has been integrated into the analysis framework *Volatility* that we will describe in depth in Section 5. Van Baar et al. (2008) have extended these works and developed a prototype that is capable of finding file remnants in RAM even if the associated process has terminated.

In spite of these features, it is important to emphasize that the VAD tree itself may be the target of a compromise attempt. An adversary with system privileges might, for instance, remove a node from the tree, similar to the DKOM attack

presented in Section 4.1, or overwrite the pointer to the `_Control_Area` structure, thereby effectively hiding a malicious artifact from the forensic practices described above (Dolan-Gavitt, 2007a; Ligh et al., 2010). Even though we are not aware of any malware species that leverage these types of modifications, sample code has been written that proves the deficiencies of many security products offered on the market to date (NT Internals, 2009). Therefore, investigators should evaluate data from various sources to obtain views from multiple angles of the compromised system.

4.6. System state- and application-specific analysis

In addition to the pieces of evidence an investigator may collect with the techniques described in the previous sections, a memory image frequently contains a myriad of information about the system state that can be of great benefit to an investigation. Especially the aforementioned `_EPROCESS` block is a source of valuable data. For instance, the `StartTime` and

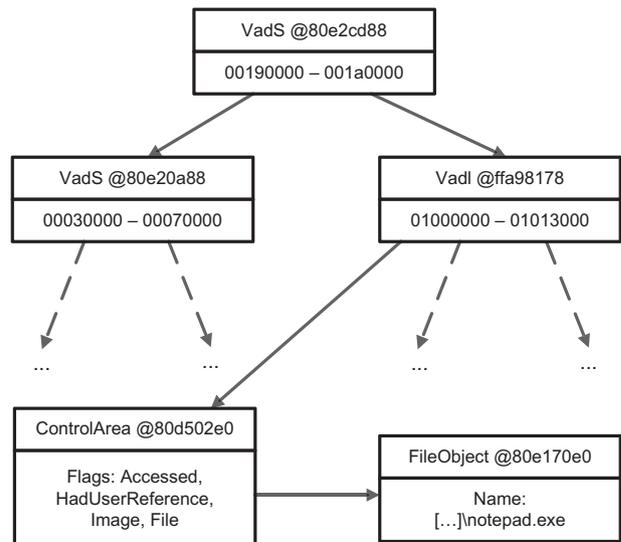


Fig. 9 – Virtual Address Descriptor (VAD) Tree (Source: Based on Dolan-Gavitt, 2007a).

¹² The root of the tree is saved in the `VADROOT` member of a process's `_EPROCESS` structure. Strategies to find this structure in memory are explained in Section 4.1.

`ExitTime` fields indicate the start and respective end time of a process and may be parsed to create forensic timelines. In addition, the periods an application has spent in system and user mode can be derived from members of the kernel process (`_KPROCESS`) block, a structure which is part of the `_EPROCESS` environment (see [Russinovich et al., 2009](#)). [Schuster \(2008b\)](#) has proven that these types of artifacts may be recovered from RAM even after a program has terminated for more than 24 h. Furthermore, with the help of the `Token` member, it is possible to reconstruct the *security context* of an application. As [Russinovich et al. \(2009, p. 473\)](#) explain, “a security context consists of information that describes the privileges, accounts, and groups associated with the process or thread”. Of particular interest is the list of user and group SIDs (*Security Identifiers*) that eventually reveal the name of the user and corresponding group account the executable was run as (see [Dolan-Gavitt, 2008d](#)).

A different research focus is set by [Stevens and Casey \(2010\)](#). They dissect the structure of the `DOSKEY` utility that is integrated into the command shell and permits editing past commands as well as displaying the command history. The latter is solely maintained in memory and is only accessible as long as the command prompt is open. As [Stevens and Casey \(2010, p. S58\)](#) point out, “[i]n practice, this makes recovering the command history difficult due to its volatile nature and the low likelihood of finding an open command window during an investigation”, even though major parts “may be recoverable from memory for some time after the windows has been closed”. The authors generate a unique signature for various `DOSKEY` elements and succeed in restoring command history objects from a number of reference data sets, including intact lists of entered commands. As these “might contain the only retrievable traces of a deleted file or suspect activity”, such types of examinations support other system state-oriented investigation methods and “can provide significant context into how and what occurred on [a] system” ([Stevens and Casey, 2010, p. S57](#)).

While most approaches we have described so far target the operating system architecture, “one of the issues currently faced in the analysis of physical memory is the recovery and use of application-level data” ([Simon and Slay, 2009, p. 996](#)). As [Simon and Slay \(2010\)](#) argue, these steps become necessary due to the increasing spread of anti-forensic technologies and a trend toward online and network applications. However, research results in this area still remain sparse to date which may be partially due to the low lifespan of data loaded in the user address space ([Solomon et al., 2007](#)). Published solutions mainly comprise instant messaging and Voice over IP (VoIP) applications ([Gao and Cao, 2010](#); [Simon and Slay, 2010](#)). Other security professionals have concentrated on finding remnants of social communication platforms in RAM ([Bryner, 2009](#); [Drinkwater, 2009](#)). With the growing use of these services, additional and more advanced techniques are likely to be developed in the future.

5. Volatility: a memory analysis framework

While most memory analysis utilities we have described in the previous sections sophisticatedly solve a specific problem,

they were typically not designed with a holistic forensic process in mind but rather as a proof-of-concept demonstration. As a consequence, many programs have their own user interface, must be invoked with different command line options, and generally neglect interprocess communication with other applications. In addition, some tools are OS-dependent and only work with specific platforms and configurations. [Vidas \(2006\)](#) notes that, for instance, the structure of the `_EPROCESS` block significantly differs across operating system versions and service pack levels, rendering hard-coded address offsets in process scanners ineffective. In sum, these characteristics force security professionals to thoroughly understand the scope and limitations of a myriad of individual solutions, and great efforts must be made to correlate results and generate homogenous reports. [Case et al. \(2008, p. S65\)](#) argue, “[a]s the complexity of systems grows rapidly, it becomes ever more difficult for an investigator to perform thorough, reliable, and timely investigations”.

In order to cope with these issues, [Walters and Petroni \(2007\)](#) suggest integrating memory forensic techniques with the digital investigation process model proposed by [Carrier and Spafford \(2003\)](#). As [Walters and Petroni \(2007, p. 2\)](#) point out, the model “allows us to organize the way we think about, discuss, and implement the procedure that are typically performed during an investigation” and “forces us to think about all phases of an investigation and how the tools and techniques that are used during a digital investigation fit together to support the process”. Their work builds the foundation for the forensic framework *Volatility* ([Volatile Systems, LLC, 2008](#)) which, in turn, has been embedded in several other application suites such as *PyFlag* ([Cohen, 2008](#)), *DFF* ([ArxSys, 2009](#)), or *PTK* ([DFlabs, 2010](#)). [Haruyama \(2009\)](#) created a collection of scripts to port *Volatility* to the commercially distributed *Encase* environment.

With respect to its architecture, *Volatility* consists of several core modules that are written in Python. The functionality of the application can be further extended with plug-ins that are provided by the developer community.¹³ Even though earlier releases were solely capable of examining images of Microsoft Windows XP, the most recent branch (1.4) also supports current operating systems, including Windows Vista and Windows 7. An overview of important program features can be seen in [Appendix A](#), the complete list of modules can be found in the work by [Ligh et al. \(2010\)](#) or when running *Volatility* with the `-h` switch.

In sum, the framework implements a great part of the concepts and methods we have outlined in this paper and has already been successfully leveraged in different scenarios ([DFRWS, 2008](#); [Smith and Cote, 2010](#)). Its architecture makes it suitable for a high degree of memory forensic-related tasks. However, it is also important to note that a high level of expertise is required to interpret extracted information and correlate results. Therefore, at the time of this writing, *Volatility* predominantly aims at academic researchers and security professionals with a strong technical background who are able to adapt the framework to their needs on a case-specific basis.

¹³ A comprehensive list of available plug-ins is maintained by members of the *Forensics Wiki* project ([Forensics Wiki, 2010](#)).

6. Conclusion and future work

The volatile storage of a computer contains a plethora of valuable information that may be key in the course of an investigation. Data found in RAM or in the system page file may significantly contribute to the reconstruction of an incident and therefore, supplement hard disk- and persistent media-oriented approaches in computer forensics. For this reason, best practices should generally include preserving those artifacts for later examination, correspondent to the guidelines as proposed by Brezinski and Killalea (2002). In sum, “during the forensic process as much attention must be paid to volatile memory as is paid to the more traditional sources of evidence” (Maclean, 2006, p. 29).

We have illustrated the prevalent concepts for creating a memory snapshot of a running machine. However, security professionals must also be aware of the scope and limitations of the different technologies. Understanding the individual benefits and drawbacks is crucial for choosing an adequate acquisition strategy. Particularly the impact of software-based solutions has been assessed only marginally at the time of this writing and must still be evaluated in more detail.

The documentation and analysis of evidence found in memory has received broad attention in research more recently. As we have outlined in Section 5, existing solutions are mainly *problem-oriented* to date though and frequently fail to provide appropriate interfaces for data import/export tasks and correlating results. The *Volatility* framework we have briefly described in the previous section addresses some of these issues and is capable of processing images in raw, crash dump, and hibernation file format. Due to its modular design, the functionality of the application can be easily extended if necessary. Case et al. (2008, p. S65) emphasize, however, that “merely swamping the user with all available data [...] falls well short of what is actually needed” and “is not, by itself, very useful”. Thus, it is also important to develop suitable visualization techniques in order to properly present the collected evidence. Only then will analysts fully recognize memory forensics as “a critical component of the digital crime scene” (Walters and Petroni, 2007, p. 15) and not as an additional burden that impedes the timely and cost-effective completion of a case.

As we have argued in the introduction of this paper, we expect the use of file and whole-disk encryption to dramatically increase in the future, both in the private as well as organizational sector. While citizens and institutions have the legitimate right to protect their data and resources on the one hand, especially law enforcement officers must be able to cope with these technologies when it becomes necessary on the other hand. We have given an overview of current approaches for recovering cryptographic keys from memory. We believe that these approaches still need to become more user-friendly and be abstracted in a way they are understandable for and applicable by technically less-sophisticated personnel.

Last but not least, with the growing number of Internet-hosted services based on virtual machines (see Section 3.3), analyzing the state of such an isolated system environment also gets more important. This so-called *virtual machine introspection* has come closer into the focus of researchers more

recently (e.g., see Dolan-Gavitt et al., 2011; Nance et al., 2009; Hay and Nance, 2008) and is a field well worth exploring from a forensic perspective.

Acknowledgments

We would like to thank Andreas Schuster for reading a previous version of this paper and giving us detailed and helpful suggestions for improvements. We would also like to thank Tilo Müller as well as Thomas Schreck for additional feedback on our work.

Table A.1 – Integral modules of the volatility memory forensics framework

| Module name | Description |
|-------------|---|
| connections | Locates the <code>_TCBTable</code> hash table in the <code>tcpip.sys</code> driver file and traverses the singly linked list of <code>_TCPT_OBJECT</code> entries to enumerate open network connections on the target system (see Section 4.4). |
| connscan2 | Scans the non-paged pool of the operating system for allocations that contain information about open network connections (see Section 4.4). |
| dlllist | Retrieves the base address, size, and path of all dynamically loaded libraries (DLLs) that are referenced by a running application. For this operation, the <code>Ldr</code> structure of the <i>Process Environment Block</i> (PEB) is parsed. It stores three doubly linked lists of <code>_LDR_DATA_TABLE_ENTRY</code> types that hold the respective information (see Section 4.5). |
| files | Retrieves the list of open file handles that is maintained by a process (see Section 4.5). |
| getsids | Reconstructs the security context of a process to retrieve the list of user and group SIDs (<i>Security Identifiers</i>) the application is associated with (see Section 4.6). |
| hivelist | Prints the virtual address and name of hive structures that internally represent parts of the system registry (see Section 4.3). |
| hivescan | Uses a signature-based approach to search for <code>_CMHIVE</code> structures in memory. These data types internally represent parts of the system registry (see Section 4.3). |
| imageinfo | Displays various meta data about the image, including the image type as well as the creation date and time. |
| kpcrscan | Scans a memory image for a <code>_KPCR</code> structure that defines the <i>Kernel Processor Control Region</i> . With the help of the structure, it is for instance possible to enumerate the list of running processes on a machine (see Section 4.1). |

Table A.1 (continued)

| Module name | Description |
|-------------|---|
| modscan2 | Scans a memory image for <code>_LDR_DATA_TABLE_ENTRY</code> objects that save the base address, size, and path of all dynamically loaded libraries (DLLs) that are referenced by a process (see Section 4.5). |
| procedump | Creates an executable (.EXE) file of a process. While the recovered file is unlikely to run due to changes in its data section, the binary can frequently be successfully disassembled for further analysis and reveal malicious activities (see Section 4.5). |
| pslist | Follows the <code>ActiveProcessLinks</code> list that is part of the <code>_EPROCESS</code> block to enumerate running processes on the system (see Section 4.1). |
| psscanner | Applies a signature-based search algorithm to locate <code>_EPROCESS</code> structures within a memory image and reveal potential <i>Direct Kernel Object Manipulation</i> (DKOM) attacks (see Section 4.1). |
| pstree | Generates a tree-like view of running system processes. The tree is derived from the output of the <code>pslist</code> command (see Section 4.1). |
| regobjkeys | Enumerates the list of open registry keys that are referenced by a process (see Section 4.3). |
| sockets | Finds the <code>_AddrObjTable</code> hash table in the <code>tcpip.sys</code> driver file and follows the singly linked list of <code>_ADDRESS_OBJECT</code> entries to enumerate listening sockets on the target system (see Section 4.4). |
| sockscan | Scans the non-paged pool of the operating system for allocations that contain information about open network sockets (see Section 4.4). |
| vaddump | Traverses the <i>Virtual Address Descriptor</i> (VAD) tree of a process and dumps the allocated memory segments to a file (see Section 4.5). |
| vadinfo | Prints detailed information about a process's <i>Virtual Address Descriptor</i> (VAD) tree that contains the range of memory addresses that are allocated by the application as well as references to loaded modules and memory-mapped files (see Section 4.5). |
| vadtrees | Creates a graphical structure of a process's <i>Virtual Address Descriptor</i> (VAD) tree (see Section 4.5). |

REFERENCES

- AccessData. Forensic toolkit 3, <http://www.accessdata.com/forensictoolkit.html>; 2010.
- Aggarwal G, Bursztein E, Jackson C, Boneh D. An analysis of private browsing modes in modern browsers. In: Proceedings of the USENIX security symposium; 2010.
- Apple Inc. Encrypting your home folder with filevault, <http://docs.info.apple.com/article.html?path=Mac/10.5/en/8736.html>; March 2010.
- Aquilina JM, Casey E, Malin CH. Malware forensics: investigating and analyzing malicious code. Syngress Publishing; 2008.
- ArxSys. Dff: digital forensics framework, <http://www.arxsys.fr/>; 2009.
- Aumaitre D. A little journey inside windows memory. Journal in Computer Virology 2009;5(2):105–17.
- Böck B. Firewire-based physical security attacks on windows 7, efs and bitlocker, http://www.securityresearch.at/publications/windows7_firewire_physical_attacks.pdf; 2009.
- Barbosa E. Finding some non-exported kernel variables in windows xp, <http://www.rootkit.com/vault/Opc0de/GetVarXP.pdf>; 2005.
- Becher M, Dornseif M, Klein CN. Firewire - all your memory are belong to us. In: Proceedings of the annual CanSecWest applied security conference; 2005.
- Beebe NL, Clark JG. Digital forensic text string searching: improving information retrieval effectiveness by thematically clustering search results. Digital Investigation 2007;4(S1):49–54.
- Beebe N, Dietrich G. A new process model for text string searching. In: Advances in digital forensics III. IFIP international federation for information processing. Boston: Springer; 2007. p. 179–91.
- Betz C. Memparser analysis tool, <http://sourceforge.net/projects/memparser/>; July 2006.
- Bilby D. Low down and dirty: anti-forensic rootkits. In: Proceedings of Ruxcon 2006; 2006.
- Boileau A. Firewire, dma & windows, <http://www.storm.net.nz/projects/16>; 2006a.
- Boileau A. Hit by a bus: physical access attacks with firewire. In: Proceedings of Ruxcon 2006; 2006b.
- Boileau A. winlockpwn, <http://www.storm.net.nz/static/files/winlockpwn>; 2008.
- Brezinski D, Killalea T. Rfc3227-guidelines for evidence collection and archiving, <http://www.faqs.org/rfcs/rfc3227.html>; February 2002.
- Brosch T, Morgenstern M. Runtime packers – the hidden problem?, <http://www.blackhat.com/presentations/bh-usa-06/BH-US-06-Morgenstern.pdf>; 2006.
- Bryner J. Facebook memory forensics, <http://computer-forensics.sans.org/blog/2009/11/20/facebook-memory-forensics>; November 2009.
- Burdach M. An introduction to windows memory forensic, http://forensic.secure.net/pdf/introduction_to_windows_memory_forensic.pdf; July 2005.
- Butler J. Fu rootkit, http://www.rootkit.com/board_project_fused.php?did=proj2012; 2005.
- Carr C. Grepxec: grepping executive objects from pool memory, <http://uninformed.org/?v=4&a=2&t=pdf>; May 2006.
- Carrier BD, Grand J. A hardware-based memory acquisition procedure for digital investigations. Digital Investigation 2004; 1(1):50–60.
- Carrier B, Spafford EH. Getting physical with the digital investigation process. International Journal of Digital Evidence 2003;2(2):1–20.
- Carvey H. The windows registry as a forensic resource. Digital Investigation 2005;2(3). 201–205.
- Carvey H. Windows forensic analysis. Syngress Publishing; 2007.
- Case A, Cristina A, Marziale L, Richard GG, Roussev V. Face: automated digital evidence discovery and correlation. Digital Investigation 2008;5(1):S65–75.
- Casey E. Digital evidence and computer crime – forensic science, computers, and the internet. 2nd ed. Academic Press; 2004.
- Chan EM, Carlyle JC, David FM, Farivar R, Campbell RH. Bootjacker: compromising computers using forced restarts. In: Proceedings of the 15th ACM conference on computer and communications security; 2008.
- Chan E, Wan W, Chaugule A, Campbell R. A framework for volatile memory forensics. In: Proceedings of the 16th ACM conference on computer and communications security; 2009.

- Chang K, Kim G, Kim K, Kim W. Initial case analysis using windows registry in computer forensics. In: Proceedings of the future generation communication and networking; 2007.
- Chow J, Pfaff B, Garfinkel T, Rosenblum M. Shredding your garbage: reducing data lifetime through secure deallocation. In: Proceedings of the 14th conference on USENIX security symposium; 2005.
- Cohen M. Pyflag – an advanced network forensic framework. In: Proceedings of the digital forensic research workshop (DFRWS); 2008.
- Computer Security Institute. 14th annual CSI computer crime and security survey; December 2009.
- Craton J. Netcat for windows, <http://joncraton.org/blog/netcat-for-windows>; 2009.
- Crazylord. Playing with windows /dev/(k)mem. Phrack Magazine 2002;59(16):1–14.
- DFRWS. Dfrws 2005 forensics challenge, <http://www.dfrws.org/2005/challenge/index.shtml>; 2005.
- DFRWS. Dfrws 2008 forensics challenge results, <http://www.dfrws.org/2008/challenge/results.shtml>; 2008.
- DFlabs. Ptk forensics, <http://ptk.dflabs.com/>; 2010.
- Darawk. Cloakdll, <http://www.darawk.com/code/CloakDll.cpp>; 2005.
- Dolan-Gavitt B. The vad tree: a process-eye view of physical memory. Digital Investigation 2007a;4(1):62–4.
- Dolan-Gavitt B. Vadtools, <http://vadtools.sourceforge.net/>; May 2007b.
- Dolan-Gavitt B. Cell index translation, <http://moyix.blogspot.com/2008/02/cell-index-translation.html>; February 2008a.
- Dolan-Gavitt B. Finding kernel global variables in windows, <http://moyix.blogspot.com/2008/04/finding-kernel-global-variables-in.html>; 2008b.
- Dolan-Gavitt B. Forensic analysis of the windows registry in memory. Digital Investigation 2008c;5(1):S26–32.
- Dolan-Gavitt B. Linking processes to users, <http://moyix.blogspot.com/2008/08/linking-processes-to-users.html>; August 2008d.
- Dolan-Gavitt B. Reading open keys, <http://moyix.blogspot.com/2008/02/reading-open-keys.html>; February 2008e.
- Dolan-Gavitt B, Srivastava A, Traynor P, Giffin J. Robust signatures for kernel data structures. In: Proceedings of the 16th ACM conference on computer and communications security; 2009.
- Dolan-Gavitt B, Payne B, Lee W. Leveraging forensic tools for virtual machine introspection. Georgia Institute of Technology; 2011. Tech. rep.
- Dornseif M, Becher M. Feuriges hacken – spaß mit firewire. In: Proceedings of the 21st chaos communication congress; 2004.
- Drinkwater R. Facebook chat forensics, <http://forensicsfromthesausagefactory.blogspot.com/2009/03/facebook-chat-forensics.html>; March 2009.
- Farmer D, Venema W. Forensic discovery. Addison Wesley; 2005.
- Forensics Wiki. List of volatility plugins, http://www.forensicswiki.org/wiki/List_of_Volatility_Plugins; February 2010.
- Foundstone, Inc. Fport, <http://www.foundstone.com/us/resources/proddesc/fport.htm>; 2000.
- Knttools with kntlist. GMG Systems, Inc, <http://gmgsystemsinc.com/knttools/>; 2007.
- Gao Y, Cao T. Memory forensics for qq from a live system. Journal of Computers 2010;5(4):541–8.
- Garcia GL. Forensic physical memory analysis: an overview of tools and techniques. Helsinki University of Technology, http://www.tml.tkk.fi/Publications/C/25/papers/Limongarcia_final.pdf; 2007. Tech. rep.
- Garner GM. Knttools with kntlist, <http://gmgsystemsinc.com/knttools/>; 2007.
- Garner GM. Forensic acquisition utilities, <http://gmgsystemsinc.com/fau/>; 2009.
- Gary HB. Fastdump – a memory acquisition tool, <https://www.hbgary.com/products-services/fastdump/>; 2009.
- Getgen K. 2009 encryption and key management industry benchmark report, <http://iss.thalesgroup.com/Resources/Webinars/Benchmarksurvey-November52009-Webinar.aspx>; 2009.
- Richard III Golden G. Scalpel: A frugal, high performance file carver, <http://www.digitalforensicsolutions.com/Scalpel/>; 2006.
- Group TC. Trusted platform module – design principles, http://www.trustedcomputinggroup.org/resources/tpm_main_specification; July 2007.
- Guidance Software. Encase enterprise platform, <http://www.guidancesoftware.com/computer-forensics-fraud-investigation-software.htm>; 2010.
- Halderman JA, Schoen SD, Heninger N, Clarkson W, Paul W, Calandrino JA, et al. Lest we remember: cold-boot attacks on encryption keys. Communications of the ACM 2009;52(5):91–8.
- Hargreaves C, Chivers H. Recovery of encryption keys from memory using a linear scan. In: Proceedings of the 2008 third international conference on availability, reliability and security; 2008.
- Haruyama T. Encase encrypt "memory forensic toolkit", <http://tharuyama.blogspot.com/2009/10/encase-encrypt-memory-forensic-toolkit.html>; October 2009. part1.
- Hay B, Nance K. Forensics examination of volatile system data using virtual introspection. ACM SIGOPS Operating Systems Review 2008;42(3):74–82.
- Hejazi S, Talhia C, Debbabi M. Extraction of forensically sensitive information from windows physical memory. Digital Investigation 2009;6(1):S121–31.
- Heninger N, Shacham H. Reconstructing rsa private keys from random key bits. Cryptology ePrint Archive 2009;(510):1–17. 2008.
- Hoglund G. The value of physical memory for incident response, <http://www.hbgary.com/wp-content/themes/blackhat/images/the-value-of-physical-memory-for-incident-response.pdf>; 2008.
- HoneyNet Project. Know your enemy – learning about security threats. Addison Wesley; 2004.
- Hulton D. Cardbus bus-mastering: Owning the laptop. In: Proceedings of ShmooCon; 2006.
- Intel Corporation. Intel® 64 and ia-32 architectures software developer's manual, <http://www.intel.com/Assets/PDF/manual/325384.pdf>; 2011.
- Internet Crime Complaint Center. 2009 ic3 annual report, http://www.ic3.gov/media/annualreport/2009_IC3Report.pdf; 2010.
- Ionescu A. Getting kernel variables from kdversionblock. part 2, <http://www.rootkit.com/newsread.php?newsid=153>; 2005.
- Kaplan, B., 2007. Ram is key – extracting disk encryption keys from volatile memory. Master's thesis, Carnegie Mellon University.
- Kdm. Ntillusion: a portable win32 userland rootkit. Phrack Magazine 2004;11(62):1–26.
- Klein T. All your private keys are belong to us – extracting rsa private keys and certificates from process memory, http://trapkit.de/research/sslkeyfinder/keyfinder_v1.0_20060205.pdf; February 2006a.
- Klein T. Process dumper, <http://www.trapkit.de/research/forensic/pd/index.html>; 2006b.
- Kornblum JD. Exploiting the rootkit paradox with windows memory analysis. International Journal of Digital Evidence 2006;5(1):1–5.
- Kornblum J. Recovering executables with windows memory analysis, <http://www.jessekornblum.com/presentations/dodcc07.pdf>; 2007a.
- Kornblum JD. Using every part of the buffalo in windows memory analysis. Digital Investigation 2007b;4(1):24–9.
- Libster E, Kornblum JD. A proposal for an integrated memory acquisition mechanism. ACM SIGOPS Operating Systems Review 2008;42(3):14–20.

- Ligh M, Adair S, Hartstein B, Richard M. Malware analyst's cookbook and DVD: tools and techniques for fighting malicious code. Wiley; 2010.
- Lynch NA, Merritt M, Weihl WE, Fekete A. Atomic transactions. Morgan Kaufmann; 1993.
- Müller T, Dewald A, Freiling F. Tresor runs encryption securely outside ram. In: Proceedings of the 20th USENIX security symposium; 2011.
- Müller T, Dewald A, Freiling FC. Aesse: A cold-boot resistant implementation of aes. In: Proceedings of the third European workshop on system security; 2010.
- Maartmann-Moe C, Thorkildsen SE, Årnes A. The persistence of memory: forensic identification and extraction of cryptographic keys. *Digital Investigation* 2009;6(S1):S132–40.
- Maclean, N.P., 2006. Acquisition and analysis of windows memory. Master's thesis, University of Strathclyde, Glasgow.
- Memory dd. ManTech CSI, Inc, <http://sourceforge.net/projects/mdd/files/>; 2009.
- Mandiant. Memoryze, http://www.mandiant.com/products/free_software/memoryze/; 2010.
- Mee V, Tryfonas T, Sutherland I. The windows registry as a forensic artefact: illustrating evidence collection for internet usage. *Digit* 2006;3(3):166–73.
- Microsoft Corporation. Notmyfault, <http://download.sysinternals.com/Files/Notmyfault.zip>; 2010 g.
- Microsoft Corporation. Debugging tools for windows, <http://www.microsoft.com/whdc/devtools/debugging/default.aspx>; 2010a.
- Microsoft Corporation. Device physicalmemory object, <http://technet.microsoft.com/de-de/library/cc78756528WS.1029.aspx>; 2010b.
- Microsoft Corporation. Exallocatepoolwithtag, <http://msdn.microsoft.com/en-us/library/ff54452028WS.1029.aspx>; 2010c.
- Microsoft Corporation. Kb244139-windows feature lets you generate a memory dump file by using the keyboard. Crash dump file generation, <http://support.microsoft.com/?20scid=kb3Ben-us3B244139&x=5&y=9>; 2010d.
- Microsoft Corporation. Kb254649-overview of memory dump file options for windows vista, windows server 2008 r2, windows server 2008, windows server 2003, windows xp, and windows 2000, <http://support.microsoft.com/?scid=kb3Ben-us3B254649&x=13&y=5>; 2010e.
- Microsoft Corporation. Kb971284-a hotfix is available to enable crashonctrlscroll support for a usb keyboard on a computer that is running windows vista sp1 or windows server 2008. Adding support for USB keyboards in Windows Vista and Server 2008 to create a crash dump file, <http://support.microsoft.com/?scid=kb3Ben-us3B244139&x=5&y=9>; 2010f.
- Microsoft Corporation. Virtuallock function, <http://msdn.microsoft.com/en-us/library/aa36689528v=vs.8529.aspx>; 2010h.
- Microsoft Corporation. Windows registry information for advanced users, <http://support.microsoft.com/kb/256986/>; February 2008.
- Microsoft Corporation. Bitlocker drive encryption, <http://technet.microsoft.com/en-us/library/cc73154928WS.1029.aspx>; February 2009.
- Miller M, Turkulainen J. Remote library injection, <http://www.nologin.org/Downloads/Papers/remote-library-injection.pdf>; June 2006.
- Moonsols. Windows memory toolkit, <http://moonsols.com/product>; 2010.
- Moore D, Paxson V, Savage S, Shannon C, Staniford S, Weaver N. Inside the slammer worm. *IEEE Security and Privacy* 2003;1(3): 33–9.
- Movall P, Nelson W, Wetzstein S. Linux physical memory analysis. In: Proceedings of the USENIX annual technical conference; 2005.
- Mozilla Foundation. Private browsing, <https://wiki.mozilla.org/PrivateBrowsing>; January 2008.
- Mrdovic S, Huseinovic A, Zajko E. Combining static and live digital forensic analysis in virtual environment. In: Proceedings of the XXII international symposium on information, communication and automation technologies; 2009.
- NT Internals. Hidden dynamic-link library detection test, http://www.ntinternals.org/dll_detection_test.php; October 2009.
- Nance K, Bishop M, Hay B. Investigating the implications of virtual machine introspection for digital forensics. In: Proceedings of the international conference on availability, reliability and security; 2009.
- Net Applications. Operating system market share, <http://marketshare.hitslink.com/operating-system-market-share.aspx?qprid=10>; December 2010.
- Okolica J, Peterson GL. Windows operating systems agnostic memory analysis. *Digital Investigation* 2010;7(1):S48–56.
- Open System Resources. Bang! – crash on demand utility, <http://www.osonline.com/article.cfm?article=153>; 2009.
- Oswald E. Toshiba's hard drive breakthrough could herald megacapacity drives, <http://www.betanews.com/article/Toshibas-hard-drive-breakthrough-could-herald-megacapacity-drives/1282246111>; August 2010.
- PGP Corporation. Pgp whole disk encryption, <http://www.pgp.com/products/wholediskencryption/index.html>; 2010.
- Panholzer P. Physical security attacks on windows vista, https://www.sec-consult.com/files/Vista_Physical_Attacks.pdf; 2008.
- Peng T, Leckie C, Ramamohanarao K. Survey of network-based defense mechanisms countering the dos and ddos problems. *ACM Computing Surveys* 2007;39(1):1–42.
- Petroni NL, Fraser T, Molina J, Arbaugh WA. Copilot – a coprocessor-based kernel runtime integrity monitor. In: Proceedings of the 13th USENIX security symposium; 2004.
- Piegdon DR, Pimenidis L. Targeting physically addressable memory. In: Proceedings of the 4th international conference on detection of intrusions and malware, and vulnerability assessment; 2007.
- Rapid7 LLC. Metasploit's meterpreter, <http://www.metasploit.com/documents/meterpreter.pdf>; 2004.
- Rare Ideas. Portableapps.com platform, <http://portableapps.com/>; 2010.
- Rolles R. Unpacking virtualization obfuscators. In: Proceedings of the 3rd USENIX conference on offensive technologies; 2009.
- Roussev V, Richard III GG. Breaking the performance wall: The case for distributed digital forensics. In: Proceedings of the digital forensic research workshop (DFRWS); 2004.
- Ruff N. Windows memory forensics. *Journal in Computer Virology* 2008;4(2):83–100.
- Ruff N, Suiche M. Enter sandman. In: Proceedings of the 5th annual PacSec applied security conference; 2007.
- Russinovich ME. Tcpview, <http://technet.microsoft.com/en-us/sysinternals/bb897437.aspx>; August 2010.
- Russinovich ME, Solomon DA, Ionescu A. Microsoft windows internals. 5th ed. Microsoft Press; June 2009.
- Rutkowska J. modgreper, <http://www.invisiblethings.org/tools/modGREPER>; June 2005.
- Rutkowska J. Beyond the cpu: defeating hardware based ram acquisition (part i: Amd case), <http://invisiblethings.org/papers/cheating-hardware-memory-acquisition-updated.ppt>; February 2007.
- Saout C. dm-crypt: A device-mapper crypto target, <http://www.saout.de/misc/dm-crypt/>; 2006.
- Savoldi A, Gubian P. Towards the virtual memory space reconstruction for windows live forensic purposes. In: Proceedings of systematic approaches to digital forensic engineering (SADFE); 2008. p. 15–22.
- Schatz B. Bodysnatcher: towards reliable volatile memory acquisition by software. *Digital Investigation* September 2007a;4:126–34.

- Schatz B. Recent developments in volatile memory forensics; November 2007b.
- Schuster A. Dmp file structure, http://computer.forensikblog.de/en/2006/03/dmp_file_structure.html; March 2006a.
- Schuster A. Pool allocations as an information source in windows memory forensics. In: Proceedings of IT-incident management & IT-forensics; 2006b. p. 15–22.
- Schuster A. Reconstructing a binary, http://computer.forensikblog.de/en/2006/04/reconstructing_a_binary.html; April 2006c.
- Schuster A. Searching for processes and threads in microsoft windows memory dumps. Digital Investigation July 2006d;3(1): 10–6, http://computer.forensikblog.de/en/2006/03/dmp_file_structure.html.
- Schuster A. 64bit crash dumps, http://computer.forensikblog.de/en/2008/02/64bit_crash_dumps.html; February 2008a.
- Schuster A. The impact of microsoft windows pool allocation strategies on memory forensics. Digital Investigation 2008b; 5(S1):S58–64.
- Scientific Working Group on Digital Evidence (SWGDE). Swgde recommended guidelines for validation testing – version 1.1, <http://www.swgde.org/documents/current-documents/2009-01-15SWGDERecommendationsforValidationTestingVersionv1.1.pdf>; 2009.
- Shamir A, van Someren N. Playing hide and seek with stored keys. In: Proceedings of the third international conference on financial cryptography; 1999.
- Sharif M, Lanzi A, Giffin J, Lee W. Automatic reverse engineering of malware emulators. In: Proceedings of the 30th IEEE symposium on security and privacy; 2009.
- Shiple TG, Reeve HR. Collecting evidence from a running computer – a technical and legal primer for the justice community, <http://www.hbgary.com/wp-content/themes/blackhat/images/collectevidenceruncomputer.pdf>; 2006.
- Simon M, Slay J. Enhancement of forensic computing investigations through memory forensic techniques. In: Proceedings of the international conference on availability, reliability and security (ARES); 2009.
- Simon M, Slay J. Recovery of skype application activity data from physical memory. In: Proceedings of the international conference on availability, reliability and security (ARES); 2010.
- Smith J, Cote M. The honeynet project forensic challenge 3: banking troubles, http://honeynet.org/challenges/2010_3_banking_troubles; May 2010.
- Smith JE, Nair R. The architecture of virtual machines. Journal of Computer 2005;38(5):32–8.
- Solomon J, Huebner E, Bem D, Szezynska M. User data persistence in physical memory. Digital Investigation 2007;4(2):68–72.
- Sophos Plc. Safeguard easy, <http://www.sophos.com/products/enterprise/encryption/safeguard-easy>; 2010.
- Sparks S, Butler J. Shadow walker – raising the bar for rootkit detection, <http://www.blackhat.com/presentations/bh-jp-05/bh-jp-05-sparks-butler.pdf>; 2005.
- Spitzner L. Definitions and value of honeypots, <http://www.spitzner.net/honeypots.html>; May 2003a.
- Spitzner L. Moving forward with definition of honeypots, <http://www.securityfocus.com/archive/119/321957/30/0/threaded>; May 2003b.
- Stevens RM, Casey E. Extracting windows command line details from physical memory. Digital Investigation 2010;7(1): S57–63.
- Stover S, Dickerson M. Using memory dumps in digital forensics. login: The USENIX Magazine 2005;30(6):43–8.
- Suiche M. Exploiting windows hibernation. In: Proceedings of Europol high tech crime expert meeting; 2008a.
- Suiche M. Sandman project, http://sandman.msuiche.net/docs/SandMan_Project.pdf; February 2008b.
- Suiche M. Advanced mac OS physical memory analysis, http://www.blackhat.com/presentations/bh-dc-10/Suiche_Matthieu/Blackhat-DC-2010-Advanced-Mac-OS-X-Physical-Memory-Analysis-wp.pdf; February 2010.
- Sutherland I, Evans J, Tryfonas T, Blyth A. Acquiring volatile operating system data tools and techniques. ACM SIGOPS Operating Systems Review 2008;42(3):65–73.
- BBN Technologies. Fred: Forensic ram extraction device, <http://www.ir.bbn.com/vkawadia/>; 2006.
- Truecrypt. TrueCrypt Foundation, <http://www.truecrypt.org/>; 2010.
- Tsow A. An improved recovery algorithm for decayed as key schedule images. Lecture Notes in Computer Science 2009; 5867:215–30.
- U.S. Department of Justice. Electronic crime scene investigation: A guide for first responders, <http://www.ncjrs.gov/pdffiles1/nij/219941.pdf>; 2008.
- U.S. Secret Service. Best Practices for seizing electronic evidence; 2006.
- United States Air Force Office of Special Investigations, March 2001. Foremost. <http://foremost.sourceforge.net/>
- VMware, Inc. What files make up a virtual machine?, http://www.vmware.com/support/ws55/doc/ws_learning_files_in_a_vm.html; 2010.
- van Baar R, Alink W, van Ballegooij A. Forensic memory analysis: files mapped in memory. In: Proceedings of the Digital Forensic Research Workshop (DFRWS); 2008.
- Vidas T. The acquisition and analysis of random access memory. Journal of Digital Forensic Practice 2006;1(4):315–23.
- Vidas T. Volatile memory acquisition via warm boot memory survivability. In: Proceedings of the 43rd Hawaii international conference on system sciences; 2010.
- Vidstrom A. Pmdump, <http://ntsecurity.nu/toolbox/pmdump/>; 2002.
- Vidstrom A. Memory dumping over firewire – uma issues, <http://ntsecurity.nu/onmymind/2006/2006-09-02.html>; 2006.
- Volatile Systems, LLC. The volatility framework: volatile memory artifact extraction utility framework, <https://www.volatilesystems.com/default/volatility>; 2008.
- Walters A, Fatkit: Detecting malicious library injection and upping the “Anti”. Tech. rep., 4 Tϕ Forensic Research. 2006.
- Walters A, Petroni NL. Volatools: integrating volatile memory forensics into the digital investigation process. In: Proceedings of Black Hat DC 2007; February 2007.
- X-Ways Software Technology AG. Winhex: computer forensics & data recovery software, hex editor & disk editor, <http://www.x-ways.net/winhex/index-m.html>; 2010.
- Young A, Yung M. Malicious cryptography – exposing cryptovirology. Wiley; 2004.
- Zhang R, Wang L, Zhang S. Windows memory analysis based on kpcr. In: Proceedings of the fifth international conference on information assurance and security; 2009.
- Zhang S, Wang L, Zhang R, Guo Q. Exploratory study on memory analysis of windows 7 operating system. In: Proceedings of the Third international conference on advanced computer theory and engineering (ICACTE); 2010.
- Zhao Q, Cao T. Collecting sensitive information from windows physical memory. Journal of Computers January 2009;4(1):3–10.