

System for Executing Encrypted Native Programs

Michael Kiperberg*, Amit Resh*, Roe Leon*, Nezer J. Zaidenberg*

*Department of Mathematical Information Technology, University of Jyväskylä, Finland
{michael, amit, roee, nezer}@trulyprotect.com

Abstract—An important aspect of protecting software from attack, theft of algorithms, or illegal software use, is eliminating the possibility of performing reverse engineering. One common method to deal with these issues is code obfuscation. However, it is proven to be ineffective. Code encryption is a much more effective means of defying reverse engineering, but it requires managing a secret key available to none but the permissible users. The authors propose a new and innovative solution. Critical functions in a protected software are encrypted using well-known encryption algorithms. Following verification by external attestation, a thin hypervisor is used as the basis of an eco-system that manages just-in-time decryption, inside the CPU, where decrypted instructions are then executed and finally discarded, while keeping the secret key and the decrypted instructions absolutely safe. The paper presents and compares two methodologies that perform the just-in-time decryption: in-place and buffered execution. The former being safer, while the latter boasts better performance

Index Terms—Trusted Computing, Hypervisor, Virtualization, Remote Attestation

I. INTRODUCTION

Digital content such as games, videos, and the like may be susceptible to unlicensed usage, which has a significant adverse impact on the profitability and commercial viability of such products. Commonly, such commercial digital content may be protected by a licensing verification program; these, however, may be circumvented by reverse engineering of the software instructions of the computer program which leaves them vulnerable to misuse.

One way of preventing circumvention of the software licensing program, may be using a method of obfuscation [1], [2]. The term obfuscation refers to making software instructions difficult for humans to understand by deliberately cluttering the code with useless, confusing pieces of additional software syntax or instructions. However, even when changing the software code and making it obfuscated, the content is still readable to the skilled hacker [3], [4].

Additionally, publishers may protect their digital content product by encryption, using a unique key to convert the software code to an unreadable format, such that only the owner of the unique key may decrypt the software code. Such protection may only be effective when the unique key is kept secured and unreachable to an adversary. Hardware based methods for keeping the unique key secured are possible [5]–[7], but may have significant deficiencies, mainly due to an investment required in dedicated hardware on the user side, making it costly, and, therefore, impractical. Furthermore, such

hardware methods have been successfully attacked by hackers [8], [9].

Software copy-protection is currently predominantly governed by methodologies based on obfuscation, which are volatile to hacking or user malicious activities. There is, therefore, a need for a better technique for protecting sensitive software sections, such as licensing code.

In this paper, we present a system that allows encrypting and executing native programs written for the x86 architecture. The system is based on the approach proposed by Averbuch et al. [10], in which an attested kernel module is responsible for decryption and execution of encrypted functions. The main deficiency of the proposed approach is the inability of the kernel module to protect itself from the operating system. As a consequence, a vulnerability in the operating system may compromise the secret key. Moreover, the attestation server has to attest not only the kernel module responsible for decryption but also the entire operating system. The complications of operating system attestation and a partial mitigation are described in [11].

This paper proposes to solve all these complications by utilizing the virtualization extension, which is available on modern processors [12], [13], in order to enable the decrypting kernel module to protect itself, thus eliminating the need for operating system attestation. Figure 1 depicts the components of the proposed system as well as their relationships. The system is deployed on three computers: a development machine, on which the program to be encrypted, is compiled and encrypted; the attestation server, which stores the decryption key, and delivers it to the target machine; and the target machine, which executes the encrypted program. A special driver, which embeds a hypervisor, is installed on the target machine prior to execution of an encrypted program. The hypervisor obtains the decryption key, which is necessary for program execution, from the attestation server, when an encrypted program is loaded to the memory.

The paper is organized as follows. Section II presents the structure of executable files, and describes the transformation applied to these files by the encryption tool. Hypervisors and their role in security are discussed in section III. Section IV outlines the attestation protocol, which is performed during hypervisor’s initialization. We discuss the execution process of an encrypted function in section V. The performance and security measurements of the system are presented in section

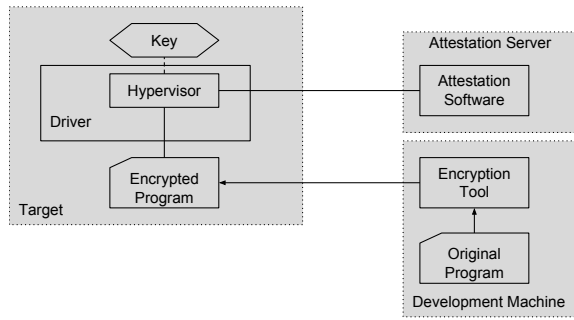


Fig. 1. Native code protection system. The original program is encrypted before its distribution. The encryption key is stored in the attestation server, which delivers it to the hypervisor in the target machine upon successful attestation. The hypervisor is initialized by a driver, which also hosts the code of the hypervisor.




Fig. 2. Structure of a Windows PE file. The structure contains a variable number of sections. Two of the most common sections are presented.

On Linux, program files, executable files and dynamic libraries, are stored in Executable And Linkable Format (ELF) format [15]. Figure 3 depicts the structure of an ELF file. An ELF file consists of a header, which is followed by data. The data may include:

- Program header table, describing zero or more segments. Only two segments can be defined as loadable: the code segment and the data segment. The code segment is loaded to memory with read-write-execute permissions, while the data segment is loaded with read-only permissions. Other segments are not loaded to memory.
- Section header table, describing zero or more sections. A typical ELF file holds a section called `.text`, which contains the code of the program.
- Data referenced by entries in the program header table or section header table.

The segments contain information that is necessary for runtime execution of the file, while the sections contain data for linking and relocation. Figure 3 depicts the structure of an ELF virtual-image at load time.

The encryption tool modifies the given PE/ELF file by introducing a new section, which stores the selected functions in encrypted form. The instructions of the original functions are partially replaced by an exception inducing instruction. We propose to use either the `halt` instruction or the `software breakpoint` instruction. The `halt` instruction is a privileged instruction, which deactivates the current processor when executed in kernel mode, but generates a general protection fault when executed in user mode. The `software breakpoint` instruction generates a breakpoint trap when executed in either kernel or user modes. Faults and traps, being types of interrupts, can be intercepted by a hypervisor, which can then decrypt and execute the original encrypted function. Another benefit of the `halt` and the `software breakpoint` instructions is that they can be represented by a single byte (0xF4 for `halt` and 0xCC for `software breakpoint`), thus allowing them to fully cover any number of bytes. The `software breakpoint` instruction is superior to the `halt` instruction in that it generates an interrupt not only in user mode but also in kernel mode.

As will be explained in section V, it is highly important to intercept control transfers that leave the encrypted function. The encryption tool disassembles the function to be encrypted and inspects its instructions. The instructions then are classified as *encryptable* and *non-encryptable*. The encryption tool classifies an instruction as non-encryptable if it might transfer control out of the encrypted function. For example, the `ret` and the `call` instructions are always classified as non-encryptable, but the `jmp` instruction is classified as non-encryptable only if its destination lays outside of the protected function's bounds or if the destination cannot be determined statically (if it is stored in a register, for instance).

The encryption tool produces two copies of the original function, the encryptable copy (EC) and the non-encryptable copy (NEC). In the EC all the non-encryptable instructions are replaced by the `halt` or the `software breakpoint` instructions. Then the encryption tool encrypts the EC and stores it in the

1 VI. Section VII discusses the limitations of the presented
 2 system and its possible extensions. Section VIII summarizes
 3 the results of this paper.

4 II. ENCRYPTION TOOL

5 The encryption tool is responsible for encryption of selected
 6 functions in a program. The user selects the functions to be
 7 encrypted by specifying their names in a configuration file.
 8 A *map file* or a *debug symbols file*, which are produced by
 9 a compiler, can then be used to translate the names of the
 10 functions to their locations in the program file.

11 On Windows, program files, executables and dynamic li-
 12 braries, are stored in Portable Executable (PE) format [14].
 13 Figure 2 depicts the structure of a PE file. The different headers
 14 define the expected location of the PE file when loaded to
 15 memory, sizes and positions of various data structures inside
 16 the PE file, the number of sections contained in this PE file,
 17 etc. The section table contains a description of each of the
 18 sections contained in the PE file. Following the section table
 19 are the sections themselves. Sections vary in their structure and
 20 purpose: the `.text` section contains the code of the program, the
 21 `.data` section contains its constants. Other sections may contain
 22 information about resources (images and sounds) embedded in
 23 the PE file or information used during exception delivery.

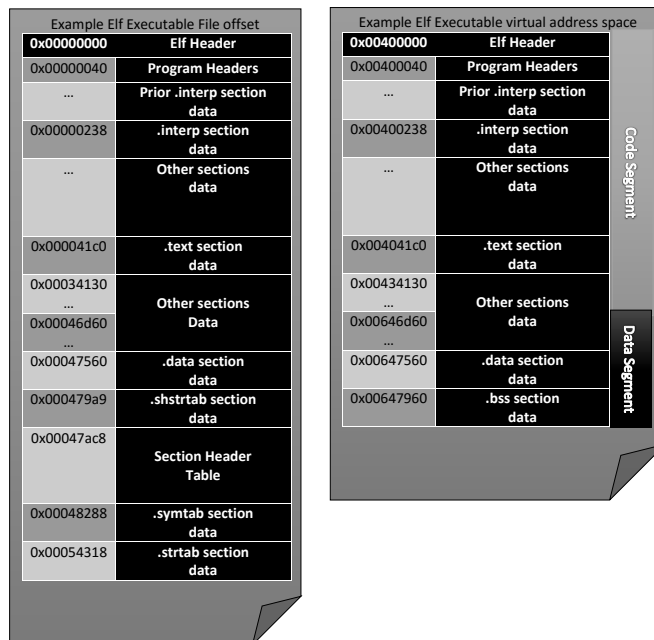


Fig. 3. The left image represents the structure of an ELF file as it is stored in disk. The right image represents the structure of an ELF file as it is loaded to memory.

1 new section. In the NEC all the encryptable instructions are
 2 replaced by the halt or the software breakpoint instructions.
 3 Then the encryption tool replaces the original function by the
 4 NEC. Figure 4 presents an example of such transformation.

III. HYPERVISOR

6 A hypervisor, also referred to as a Virtual Machine Monitor
 7 (VMM), is software, which may be hardware-assisted, to
 8 manage multiple virtual machines on a single system [16]. The
 9 hypervisor virtualizes the hardware environment in a way that
 10 allows several virtual machines, running under its supervision,
 11 to operate in parallel over the same physical hardware plat-
 12 form, without obstructing or impeding each other. Each virtual
 13 machine has the illusion that it is running unaccompanied on
 14 the entire hardware platform. The hypervisor is referred to as
 15 the *host*, while the virtual machines are referred to as *guests*.

16 A virtual machine control structure (VMCS) is defined
 17 for each virtual environment managed by a virtual machine
 18 monitor (VMM) [12]. This structure defines the values of
 19 privileged registers, the location of the interrupt descriptors
 20 table, and additional values that constitute the internal state
 21 of the virtual environment. In addition, this structure defines
 22 the events that the VMM is configured to intercept, and the
 23 address of the function that should handle the interception.
 24 The act of control transfer from the virtual environment to a
 25 predefined function is called *vm-exit* and the act of control
 26 transfer from the function back to the virtual environment
 27 is called *vm-entry*. Upon *vm-exit* the function can determine

the reason of the *vm-exit* by examining the fields of the
 VMCS and altering them, thus altering the state of the virtual
 environment as it wishes. Finally, the VMCS can define a
 mapping between the physical memory as it is perceived
 by the virtual environment and the actual physical memory.
 As a consequence, the VMM can prevent access to some
 physical pages by the virtual environment. Moreover, the
 virtual environment will be unaware of this situation.

We propose to use a hypervisor for securing a single
 guest. Rather than wholly virtualizing the hardware platform,
 a special breed of hypervisor, called a *thin hypervisor*, is
 used [17], [18]. A thin hypervisor is configured to intercept
 only a small portion of events. All other events are processed
 without interception, directly, by the OS. A thin hypervisor
 only intercepts the set of events that allows it to protect an
 internal secret (such as a cryptographic key) and protect itself
 from subversion. Figure 5 depicts a thin hypervisor supporting
 a single guest. Since a thin hypervisor does not control most
 of the OS interaction with the hardware, multiple OS are not
 supported. On the other hand, system performance is kept at
 an optimum.

A thin hypervisor facilitates a secure environment by: (a)
 setting aside portions of memory that cannot be accessed
 by the guest, (b) storing the cryptographic key in privileged
 registers, and (c) intercepting privileged instructions that may
 compromise its protected memory or the cryptographic key.

Once this environment is correctly configured, a thin hy-

28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

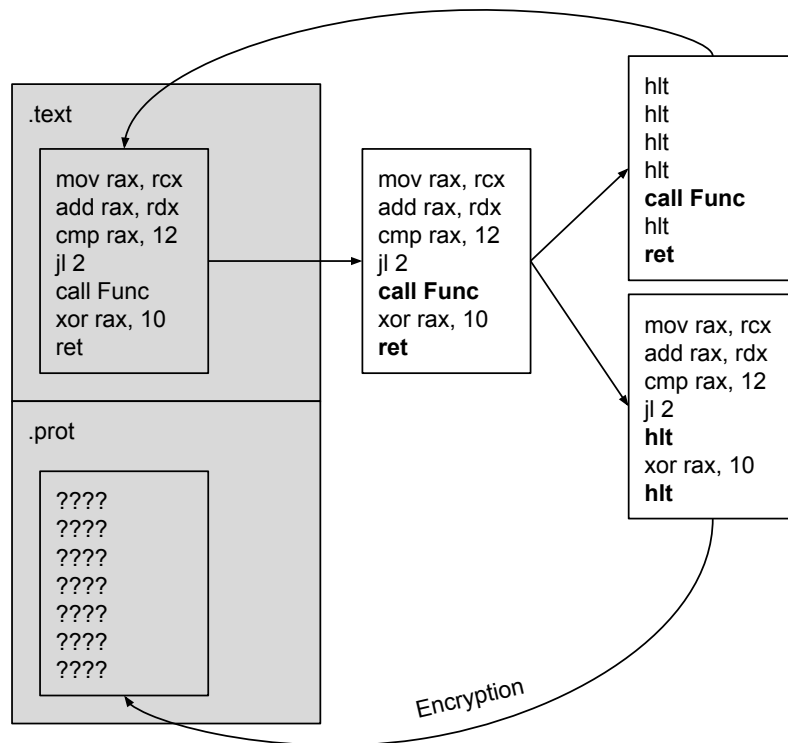


Fig. 4. Example of an encryption process of a single function. The encryption begins by classifying instruction is encryptable (normal face) and non-encryptable (bold face), and creating to copies. The complementary instructions in each copy are replaced by halts. Finally, one copy is written over the original functions, and the other is encrypted and added to the special section.

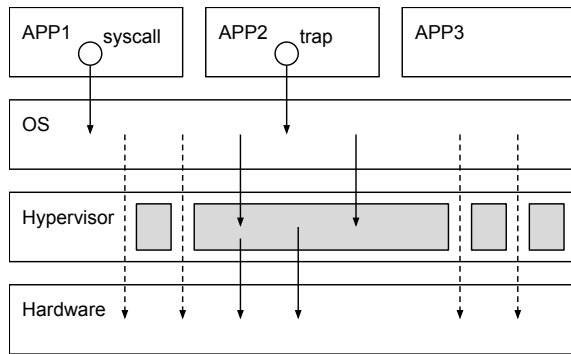


Fig. 5. Thin hypervisor. The hypervisor runs in a higher privilege level than the operating system. System calls, traps, exceptions, and other interrupts, transfer control from user mode applications to the operating system. The operating system handles these conditions by requesting some service from the underlying hardware. A thin hypervisor can intercept some of those requests and handle them according to some policy.

1 pervisor can be utilized to carry out specific operations, which
 2 may include use of the cryptographic key, in a protected region
 3 of memory. As a result of the tightly configured intercepts and
 4 absolute control of the protected memory regions, this activity
 5 can be guaranteed to protect both the cryptographic key and
 6 the operations results.

IV. REMOTE ATTESTATION

7
 8 The problem of remote software authentication, determining
 9 whether a remote computer system is running the correct
 10 version of a software, is well known [5], [19]–[25]. Equipped
 11 with a remote authentication method, a service provider can
 12 prevent an unauthenticated remote software from obtaining
 13 some secret information or some privileged service. For exam-
 14 ple, only authenticated gaming consoles can be allowed to
 15 connect to the gaming networks [26]–[28], and only authenti-
 16 cated bank terminals can be allowed to fetch records from the
 17 bank database [29].

18 The research in this area can be divided into two ma-
 19 jor branches: hardware assisted authentication [5]–[7] and
 20 software-only authentication [19]–[22]. In this paper we con-
 21 centrate on software-only authentication, although the system
 22 can be adapted to other authentication methods, as well.
 23 The authentication entails simultaneously authenticating some
 24 software component(s) or memory region, as well as verifying
 25 that the remote machine is not running in virtual or emulation
 26 mode. Software-only authentication methods may also involve
 27 a challenge code, that is sent by the authentication authority,
 28 and executed on the remote system. The challenge code
 29 computes a result that is then transmitted back to the authority.
 30 The authority deems the entity to be authenticated if the
 31 result is correct and was received within a predefined time-




Fig. 6. The attestation protocol between the authentication authority and the target machine. The protocol consists of four messages. The first two messages are sent unencrypted, while the two last messages are encrypted. The third message is encrypted by the public key of the authentication authority and the fourth message is encrypted by the random value transmitted in the third message.

a proof of target’s authenticity. The authentication authority encrypts the secret information with a random value obtained from message (3) acting as the encryption key, and transmits the encrypted message to the target machine.

V. ENCRYPTED INSTRUCTIONS EXECUTION

In order to execute an encrypted program, the user must first install the driver, which encapsulates the hypervisor. The driver monitors the PE files loaded by the OS, and keeps track of PE files that contain the special encrypted functions section. When the first such PE file is loaded, the driver initializes the hypervisor. During the initialization, the driver communicates with the authentication authority, passes the attestation verification, obtains the cryptographic key, and enters a virtualized state.

The hypervisor is configured to intercept the general protection fault. When a protected program transfers control to an encrypted function, the processor attempts to execute the halt instruction, which induces a general protection fault, thus transferring control to the hypervisor. General protection faults rarely occur during the normal course of program execution, since they usually cause the program to terminate abruptly. Nevertheless, the hypervisor uses the data structures prepared by the encryption tool to test whether the general protection fault occurred during execution of an encrypted function.

The hypervisor injects the interrupt back to the guest, if it was not caused by an encrypted function execution. Otherwise, the hypervisor decrypts the function and starts its execution. Since during its execution, the function is stored in memory in unencrypted form, it is highly important to ensure that no other code has access to the decrypted instructions of the function. We note that in modern processors, several execution units (logical processors) can execute programs concurrently. Therefore, we must ensure that programs executed by all execution units have no access to the unencrypted instructions.

We present two approaches to sensitive functions execution: *in-place execution* and *buffered execution*.

A. In-place Execution

According to this approach the hypervisor can be in one of two states: cold or hot. In the cold state the memory does not contain any sensitive information and only the cryptographic key and the hypervisor’s state must be protected. This is the regular mode of operation described in section III. The hypervisor switches to the hot state when the memory contains sensitive information, which cannot be protected by a regular memory protection technique (using EPT), since its physical location is not known (or not constant). This switch occurs when the hypervisor triggers execution of a decrypted function.

In the following description, we assume that the encryption tool uses halt as a replacement instruction, but the same is true when the software breakpoint instruction is used.

At initialization the hypervisor’s state is set to cold. In this state, in addition to the regular protection means described in section III, the hypervisor intercepts general protection faults.

1 frame. The underlying assumption, which is shared by all such
 2 authentication methods, is that only an authentic system can
 3 compute the correct result within the predefined time-frame.
 4 The methods differ in the means by which (and if) they satisfy
 5 this underlying assumption.

6 Kennell and Jamieson proposed [19] a method that produces
 7 the result by computing a cryptographic hash of a specified
 8 memory region. Any computation on a complex instruction set
 9 architecture (Pentium in this case) produces side effects. These
 10 side effects are incorporated into the result after each iteration
 11 of the hashing function. Therefore, an adversary, trying to
 12 compute the correct result on a non-authentic system, would
 13 be forced to build a complete emulator for the instruction
 14 set architecture to compute the correct side effects of every
 15 instruction. Since such an emulator performs tens and hun-
 16 dreds of native instructions for every simulated instruction,
 17 Kennell and Jamieson conclude that it will not be able to
 18 compute the correct result within the predefined time-frame.
 19 The method of Kennell and Jamieson was further adapted to
 20 modern processors [30]. The adaptation solves the security is-
 21 sues that arise from the availability of virtualization extensions
 22 and multiplicity of execution units.

23 The authentication protocol is depicted in Figure 6. The
 24 initial messages of the protocol carry information about the
 25 current configuration of the target machine. Following this
 26 exchange, the authentication authority transmits a message
 27 containing the challenge code to be executed on the target
 28 machine. The target machine executes the challenge, which
 29 computes a result, that is a cryptographic hash of some
 30 memory region, possibly with some additional information.
 31 The target machine, concatenates a randomly generated num-
 32 ber to the result, encrypts both values with the public key
 33 of the authentication authority, and transmits the encrypted
 34 message. The authentication authority verifies that the result
 35 is correct and was received within a predefined time-frame. If
 36 both are true the target machine is considered authentic. The
 37 authentication authority then shares some secret information
 38 with the target machine. This secret information constitutes

39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92

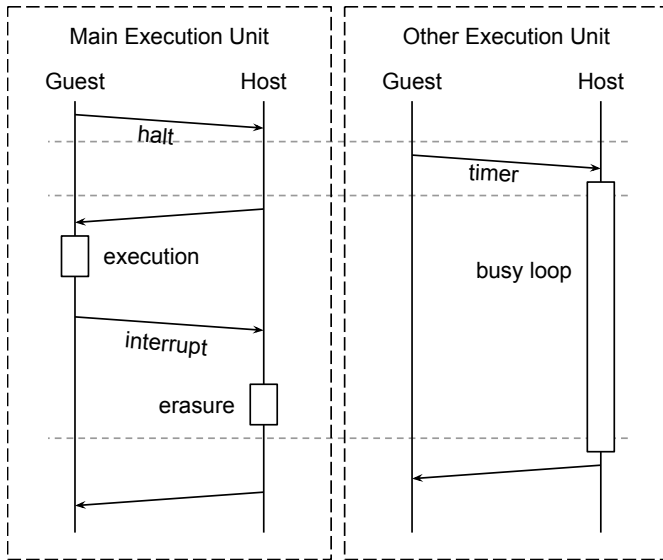


Fig. 7. Example of encrypted function execution. The figure depicts two execution units, each with two alternating states: guest and host. The dashed horizontal lines are synchronization barriers, i.e. everything above the line is guaranteed to complete before anything below the line starts.

with a specially crafted IDT. In this special IDT each handler induces a vm-exit, for example, by executing the CPUID instruction. The hypervisor intercepts this instruction, realizes that an interrupt at vector N occurred and switches to cold mode. The hypervisor proceeds by installing the original IDT and moves the guest's instruction pointer to point to the N th interrupt handler of the original IDT.

B. Buffered Execution

In the following description, we assume that the encryption tool uses halt as a replacement instruction for NECs and software breakpoint as a replacement instruction for ECs.

According to this approach, the hypervisor has only one state, in which it protects itself as described in section III. In addition, the hypervisor configures itself to intercept general protection faults. Execution of halt instructions induces a general protection fault, which causes a vm-exit and transfers control to the hypervisor. The hypervisor inspects the source of the general protection fault, and fetches the EC that corresponds to this NEC.

When the EC is resolved, the hypervisor decrypts it into a pre-allocated memory buffer, which is protected by the hypervisor. The decrypted EC will be executed in host mode, thus allowing it to reside in an EPT-protected buffer. Since the decrypted instructions are inaccessible by any other execution unit (in guest mode), there is no need to suspend them. Likewise, since the encrypted instructions are executed inside the hypervisor, there is no need to modify the IDT of the guest. Finally, there is no need to perform the costly transitions to and from the guest after every decryption. All these improve the overall performance of the system by a large factor.

The x86 instruction set architecture defines many memory access instructions as *relative*, meaning that their arguments should not be interpreted as actual memory locations but rather they should be interpreted as offsets from the current value of the instruction pointer. As a consequence, the same instruction may have different interpretations when executed at different locations. Therefore we must execute the decrypted EC at its natural location. In order to achieve this, the hypervisor modifies the virtual page table of the current process by mapping the virtual page containing the NEC to the physical address of the pre-allocated buffer containing the decrypted EC. Figure 8 depicts this transformation.

The control flow during the execution of an encrypted function is illustrated in Figure 9. The process begins when an encrypted function is called. The first instruction in the NEC is the halt instruction; its execution triggers the general protection exception, which induces a vm-exit. The hypervisor prepares the system for buffered execution by performing the following steps: (1) the EC is decrypted into a pre-allocated buffer; (2) the virtual page table is modified to map the natural location of the function to the pre-allocated buffer, as illustrated in Figure 8; (3) the values of the guest registers, which were stored during the vm-exit transition, are restored; (4) the decrypted function is called. The decrypted function executes until an interrupt occurs. The interrupt can

1 An encrypted function, which was overwritten by the NEC
 2 consists mainly of halt instructions. Execution of any of these
 3 instructions induces a general protection fault, which causes a
 4 vm-exit and transfers control to the hypervisor. The hypervisor
 5 inspects the source of the general protection fault, and fetches
 6 the EC that corresponds to this NEC. Then the hypervisor
 7 switches to hot mode and decrypts the EC into its natural
 8 location, currently occupied by the NEC (the NEC is saved in
 9 a different location for future use).

10 During the switch to hot mode, the hypervisor freezes
 11 all other execution units, and configures itself to intercept
 12 all interrupts. This behaviour guarantees that the function in
 13 its decrypted form cannot be read by any other, potentially
 14 malicious, code, simply because no other code can run in hot
 15 mode. We note that all the control transfer instructions in the
 16 EC are replaced by the halt instruction, which induces a vm-
 17 exit.

18 When a vm-exit occurs in hot mode, the hypervisor first
 19 replaces the decrypted function with the NEC, and switches
 20 to cold mode. Following this, the hypervisor resumes all the
 21 execution units, configures itself to intercept only general
 22 protection faults, and returns control to the guest. Figure 7
 23 depicts the control flow during encrypted function execution.

24 We suggest to freeze other execution units by inducing
 25 a vm-exit on each execution unit, and running a busy loop
 26 until the hypervisor switches back to cold mode. A vm-exit
 27 can be induced either implicitly with a timer or explicitly by
 28 sending an inter-processor interrupt (IPI). The former solution
 29 is much easier to implement but the later solution is much
 30 more efficient.

31 The hypervisor intercepts interrupts in hot mode by replac-
 32 ing the original interrupt descriptor table (IDT) of the OS

33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87

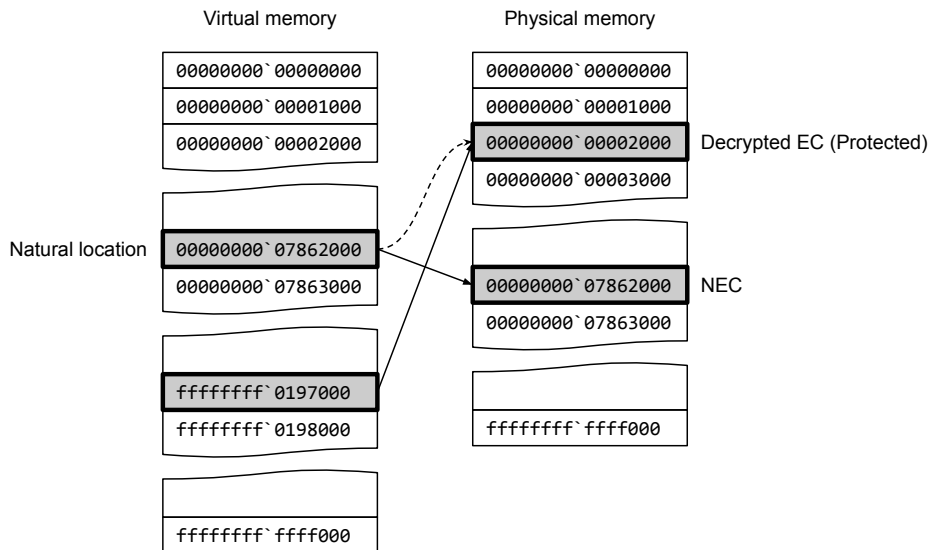


Fig. 8. Memory layout during buffered execution. The functions resided at virtual address 7862000, which is mapped to the physical address 7862000 (a coincidence). The encrypted code is decrypted to virtual address ffffffff`0197000 which is mapped to the physical address 2000. The hypervisor changes the mapping of the virtual address 7862000 to map the physical address 2000.

1 be triggered by a software breakpoint instruction or by some
 2 other condition, e.g., a page fault. In both cases the hypervisor
 3 suspends the buffered execution by performing the following
 4 steps: (1) the values of the registers are stored to a memory
 5 region from which they will be restored during vm-entry; (2)
 6 the virtual page table is restored to its original state; (3) the
 7 decrypted EC is erased. If the interrupt was triggered by a
 8 software breakpoint instruction, the hypervisor resumes the
 9 guest immediately. However, if the interrupt was triggered by
 10 some other condition, the hypervisor injects the interrupt to the
 11 guest, and then resumes it. The interrupt injection mechanism
 12 allows the hypervisor to delegate the responsibility of interrupt
 13 handling to the operating system. Figure 9 illustrates the
 14 simple case of software breakpoint interrupt.

15 This approach is more efficient but potentially less secure
 16 than the in-place execution. According to this approach, the
 17 decrypted functions are executed inside the hypervisor itself.
 18 As a consequence these functions have the same privileges as
 19 the hypervisor. In particular, they can read and write memory,
 20 which is otherwise inaccessible to any code external to the
 21 hypervisor. One can argue that it is impossible for an adversary
 22 to replace the EC with random code, without knowing the
 23 cryptographic key. However unfortunately, it is possible that
 24 some memory manipulation can be performed indirectly by
 25 modifying the data on which the encrypted function works.
 26 Nevertheless, although possible, it seems to be extremely
 27 difficult to manipulate the behaviour of unknown code through
 28 its data.

VI. MEASUREMENTS

30 This section presents a performance analysis of the two
 31 execution methods that were described in section V. The

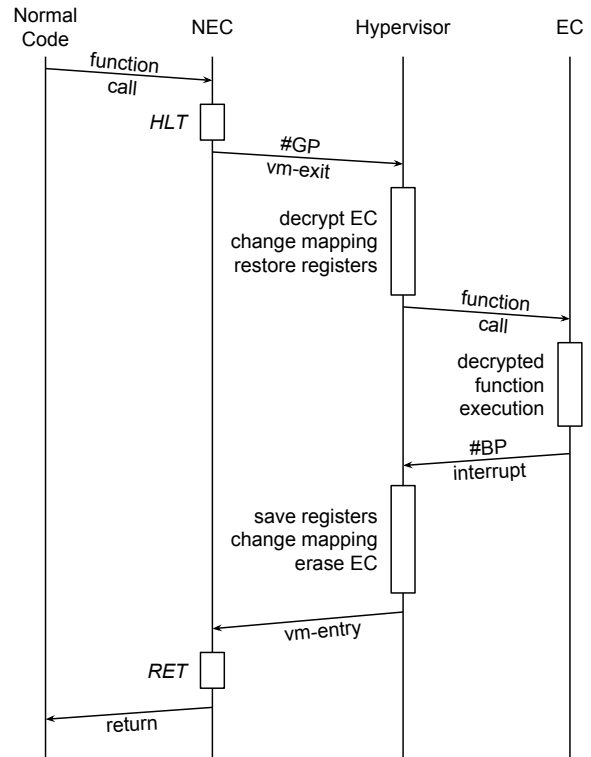


Fig. 9. Example of encrypted function execution in buffered execution mode. The figure depicts the control flow during the execution of an encrypted function.

measurements were performed on programs with a single




Fig. 10. Execution modes. The left column represents the guest mode, while the right column represents the host mode. The lower row represents the kernel mode, while the upper row represents the user mode. The host mode can protect itself from the guest mode through the EPT mechanism. The kernel mode can protect itself from the user mode through the virtual memory protection mechanism.

- vm-exit 42
- erasure of the decrypted function 43
- execution units unfreezing 44
- vm-entry 45

Figure 12 presents the execution time of each step. The buffered execution method consists of the following steps:

- vm-exit 48
- virtual page table modification 49
- decryption 50
- instructions execution 51
- erasure of the decrypted function and virtual page table restoration 52
- vm-entry 53

Figure 13 presents the execution time of each step.

Finally, figure 14 compares the execution times of the in-place and the buffered execution methods for functions of different lengths.

VII. FUTURE WORK

As was explained above, the buffered execution method is superior to the in-place execution method in terms of performance. Unfortunately, the buffered execution method allows an adversary to access regions of memory that are normally protected by the hypervisor. Consider the *memcpy* function, for example. Assume that this function is encrypted and is now being executed by the hypervisor in buffered execution mode. By specifying the address of the VMCS structure in the *source* or *destination* argument, an adversary can inspect and modify the control structures of the hypervisor. Moreover, since the hypervisor executes in kernel mode, the protected function can access OS memory region and execute privileged instructions.

Fortunately, the x86 instruction set architecture provides a great variety of memory protection mechanisms, which can be utilized by the buffered execution method. One such mechanism is the virtual memory protection, which is available in both 32- and 64-bit execution modes. The virtual memory protected mechanism allows to specify a separate set of accessibility rights for kernel mode and user mode. Similarly, the hypervisor’s memory protection (virtualization, to be precise) mechanism, called the Extended Page Table (EPT) on Intel processors, allows to specify a separate set of accessibility rights for host mode and guest mode. The different modes of execution and the protection mechanisms are summarized in Figure 10.

The in-place execution method utilizes the EPT to protect hypervisor’s control structures and other sensitive data from an adversary. We propose to use the virtual memory protection mechanism in the buffered execution method. In particular, the buffered execution method can execute the decrypted function in user mode inside the host mode (the upper right block in Figure 10); this mode is not used by the system described in this paper. In this mode we can prevent attempts to execute privileged instructions or access the hypervisor’s control structures.

1 encrypted function. The functions that were chosen do not call
 2 other functions. Therefore, these functions can be executed at
 3 once. The size of the functions varied, as will be explained
 4 below.

5 According to the in-place execution method, at the begin-
 6 ning of each execution cycle, the main execution unit freezes
 7 other execution units. This is usually accomplished with the
 8 following sequence of actions:

- 9 1) The main execution unit writes *freeze requests* to the
 10 data structures of other execution units.
- 11 2) The main execution unit triggers a vm-exit in other
 12 execution units.
- 13 3) Other execution units inspect their data structures.
- 14 4) Other execution units enter a busy-loop.
- 15 5) The main execution unit proceeds (to decryption, exe-
 16 cution, etc.).

17 Usually, in order to perform some operation on a different
 18 execution unit, the current execution unit sends an inter-
 19 processor interrupt (IPI), which triggers an interrupt service
 20 routine. This is probably the most efficient, and at the same
 21 time, the most complicated solution. We chose to induce
 22 a vm-exit periodically on every execution unit, through the
 23 *preemption-timer* field of the VMCS, which defines the maxi-
 24 mal amount of time that the guest can execute before returning
 25 to the hypervisor. Figure 11 depicts the relation between the
 26 execution time of a function and the preemption-timer interval.
 27 The figure also includes the execution times of the function
 28 under the following conditions:

- 29 • Regular, unencrypted, execution time.
- 30 • Execution time using the buffered method.
- 31 • Execution time using the in-place method, which skips
 32 the freeze step.

33 Figures 12 and 13 present the time division between the
 34 steps of the in-place and the buffered execution methods,
 35 respectively. The in-place execution method consists of the
 36 following steps:

- 37 • vm-exit
- 38 • execution units freezing
- 39 • decryption
- 40 • vm-entry
- 41 • instructions execution




Fig. 11. Comparison of execution times as a function of the preemption-timer value. The columns represent the following cases (from left to right): regular execution — no encryption, buffered execution, in-place execution without freezing, in-place execution in which the preemption timer is set to 16, in-place execution in which the preemption timer is set to 32, in-place execution in which the preemption timer is set to 64.




Fig. 12. Execution times of the steps in the in-place execution method (in which the preemption timer was set to 16). The columns represent the following steps (from left to right): entering and exiting the host; decrypting the function; executing the function; entering and exiting the host; erasing the decrypted function.




Fig. 13. Execution times of the buffered execution method steps. The columns represent the following steps (from left to right): entering and exiting the host; modifying the virtual page table and restoring the guest register; decrypting the function; executing the function; erasing the decrypted function, saving the registers, and restoring the virtual page table.




Fig. 14. Execution times of the in-place and the buffered execution methods for functions of different lengths. The small function is 92 bytes long. The medium function is 968 bytes long. The large function is 5056 bytes long.

1 The hypervisor can transit to this mode by executing the
 2 iret instruction, which is usually used to terminate an interrupt
 3 handler. This instruction modifies the execution location and
 4 the execution mode (from kernel to user). Since the execution
 5 takes place in host mode, interrupts cannot be intercepted
 6 by the hypervisor through configuration of the VMCS. The
 7 hypervisor is forced to use the IDT, which allows the kernel
 8 to specify the interrupt service routines for each of the 256
 9 interrupt vectors. Upon interrupt, the interrupt service routine
 10 can decide whether to handle the interrupt inside the hypervisor
 11 or inject it to the guest.

12 We believe that the described approach will substantially
 13 improve the security of the buffered execution method, thus
 14 making it absolutely superior to in-place execution.

VIII. CONCLUSIONS

15
 16 We present research pertaining to the methodologies of ex-
 17 ecuting encrypted native machine-code, where decryption and
 18 execution are done on the fly and secure with a thin hypervisor.
 19 Two alternative methods are considered: *in-place* and *buffered*
 20 — that trade security for performance. The in-pace method
 21 executes decrypted-code in guest mode, thereby limiting the
 22 functionality of the decrypted function to whatever a guest
 23 may perform. In buffered execution method, the decrypted
 24 function executes in host mode, penitentially incurring the risk
 25 of a rogue implementation accessing sensitive memory areas.
 26 For this reason the in-place method is considered safer. How-
 27 ever, in modern multi-processor systems, the in-place method
 28 requires controlling (freezing) other execution units, while a
 29 single execution unit executes decrypted code. This requires
 30 larger overhead when compared to the buffered method and
 31 thus has a performance toll. Measurements show that the larger
 32 overhead is more significant for larger functions. The reason
 33 for this is related to the fact that overhead is acquired during
 34 transitions between cold to hot and hot to cold modes in the
 35 in-place method, as compared to transitions between host-
 36 execution of decrypted code and guest-execution of interrupts.

Larger functions acquire more transitions, therefore overhead
 is more prominent in the in-place method. Given these results
 our conclusions are to use the (safer) in-place methodology for
 short functions (smaller than 1000 bytes). For medium (larger
 than 1000 bytes), allow a user-defined switch in the encryption
 tool to prefer security, in which case in-place shall be used, or
 performance, in which case buffered shall be used. In future
 work we plan to augment the buffered method to overcome
 its potential security flaws and render it the single and best
 alternative to use.

REFERENCES

- [1] *Themida*, <http://www.oreans.com/>, Oreans.
- [2] *VMPprotect*, <http://vmpsoft.com/>, VMPprotect Software.
- [3] R. Rolles, “Unpacking Virtualization Obfuscators,” in *Proceedings of the 3rd USENIX Conference on Offensive Technologies*, ser. WOOT’09. Berkeley, CA, USA: USENIX Association, 2009, pp. 1–1. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1855876.1855877>
- [4] L. Bohne, “Pandora’s Bochs: Automated Unpacking of Malware,” 2008.
- [5] D. Schellekens, B. Wyseur, and B. Preneel, “Remote Attestation on Legacy Operating Systems with Trusted Platform Modules,” *Sci. Comput. Program.*, vol. 74, no. 1-2, pp. 13–22, Dec. 2008. [Online]. Available: <http://dx.doi.org/10.1016/j.scico.2008.09.005>
- [6] S. Pearson, *Trusted Computing Platforms: T CPA Technology in Context*. Upper Saddle River, NJ, USA: Prentice Hall PTR, 2002.
- [7] P. England, B. Lampson, J. Manferdelli, M. Peinado, and B. Willman, “A Trusted Open Platform,” *Computer*, vol. 36, no. 7, pp. 55–62, Jul. 2003. [Online]. Available: <http://dx.doi.org/10.1109/MC.2003.1212691>
- [8] C. Tarnovsky, “Semiconductor Security Awareness Today and yesterday,” in *Blackhat*, 2010. [Online]. Available: <https://www.youtube.com/watch?v=WXX00tRKOlw>
- [9] —, “Attacking TPM part two,” in *Defcon*, 2012. [Online]. Available: https://www.youtube.com/watch?v=Ed_9p7E4jIE
- [10] A. Averbuch, M. Kiperberg, and N. J. Zaidenberg, “Truly-Protect: An Efficient VM-Based Software Protection,” *Systems Journal, IEEE*, vol. 7, no. 3, pp. 455–466, 2013.
- [11] M. Kiperberg and N. J. Zaidenberg, “Efficient Remote Authentication,” in *The Journal of Information Warfare*, vol. 12, no. 3, 2013.
- [12] *Intel 64 and IA-32 Architectures Software Developer’s Manual - Volume 3*, Intel Corporation, August 2007.
- [13] “AMD64 Architecture Programmer’s Manual Volume 2: System Programming,” AMD, 2010.
- [14] M. Pietrek, “An in-depth look into the Win32 portable executable file format,” in *MSDN Mag.* 17, 2, 2002, pp. 80–90.

- 1 [15] E. Youngdale, "Kernel korner: The elf object file format by dissection,"
2 *Linux Journal*, vol. 1995, no. 13es, p. 15, 1995.
- 3 [16] G. J. Popek and R. P. Goldberg, "Formal Requirements for Virtualizable
4 Third Generation Architectures," *Commun. ACM*, vol. 17, no. 7, pp.
5 412–421, Jul. 1974. [Online]. Available: [http://doi.acm.org/10.1145/](http://doi.acm.org/10.1145/361011.361073)
6 361011.361073
- 7 [17] T. Shinagawa, H. Eiraku, K. Tanimoto, K. Omote, S. Hasegawa,
8 T. Horie, M. Hirano, K. Kourai, Y. Oyama, E. Kawai, K. Kono, S. Chiba,
9 Y. Shinjo, and K. Kato, "Bitvisor: A thin hypervisor for enforcing i/o
10 device security," in *Proceedings of the 2009 ACM SIGPLAN/SIGOPS*
11 *International Conference on Virtual Execution Environments*, ser. VEE
12 '09. New York, NY, USA: ACM, 2009, pp. 121–130. [Online].
13 Available: <http://doi.acm.org/10.1145/1508293.1508311>
- 14 [18] Y. Chubachi, T. Shinagawa, and K. Kato, "Hypervisor-based Prevention
15 of Persistent Rootkits," in *Proceedings of the 2010 ACM Symposium on*
16 *Applied Computing*, ser. SAC '10. New York, NY, USA: ACM, 2010,
17 pp. 214–220. [Online]. Available: [http://doi.acm.org/10.1145/1774088.](http://doi.acm.org/10.1145/1774088.1774131)
18 1774131
- 19 [19] R. Kennell and L. H. Jamieson, "Establishing the Genuinity of Remote
20 Computer Systems," in *Proceedings of the 12th Conference on USENIX*
21 *Security Symposium - Volume 12*, ser. SSYM'03. Berkeley, CA,
22 USA: USENIX Association, 2003, pp. 21–21. [Online]. Available:
23 <http://dl.acm.org/citation.cfm?id=1251353.1251374>
- 24 [20] A. Seshadri, M. Luk, E. Shi, A. Perrig, L. van Doorn, and P. Khosla,
25 "Pioneer: Verifying code integrity and enforcing untampered code
26 execution on legacy systems," in *Proceedings of the Twentieth*
27 *ACM Symposium on Operating Systems Principles*, ser. SOSP '05.
28 New York, NY, USA: ACM, 2005, pp. 1–16. [Online]. Available:
29 <http://doi.acm.org/10.1145/1095810.1095812>
- 30 [21] Q. Yan, J. Han, Y. Li, R. H. Deng, and T. Li, "A software-based root-of-
31 trust primitive on multicore platforms," in *Proceedings of the 6th ACM*
32 *Symposium on Information, Computer and Communications Security*,
33 ser. ASIACCS '11. New York, NY, USA: ACM, 2011, pp. 334–343.
34 [Online]. Available: <http://doi.acm.org/10.1145/1966913.1966957>
- 35 [22] A. Seshadri, A. Perrig, L. van Doorn, and P. Khosla, "SWATT: softWare-
36 based attestation for embedded devices," in *Security and Privacy, 2004.*
37 *Proceedings. 2004 IEEE Symposium on*, May 2004, pp. 272–282.
- 38 [23] C. Castelluccia, A. Francillon, D. Perito, and C. Soriente, "On
39 the Difficulty of Software-based Attestation of Embedded Devices,"
40 in *Proceedings of the 16th ACM Conference on Computer and*
41 *Communications Security*, ser. CCS '09. New York, NY, USA: ACM,
42 2009, pp. 400–409. [Online]. Available: [http://doi.acm.org/10.1145/](http://doi.acm.org/10.1145/1653662.1653711)
43 1653662.1653711
- 44 [24] A. Seshadri, M. Luk, A. Perrig, L. van Doorn, and P. Khosla, "Scuba:
45 Secure code update by attestation in sensor networks," in *Proceedings*
46 *of the 5th ACM Workshop on Wireless Security*, ser. WiSe '06.
47 New York, NY, USA: ACM, 2006, pp. 85–94. [Online]. Available:
48 <http://doi.acm.org/10.1145/1161289.1161306>
- 49 [25] Y. Yang, X. Wang, S. Zhu, and G. Cao, "Distributed software-
50 based attestation for node compromise detection in sensor networks,"
51 in *Proceedings of the 26th IEEE International Symposium on*
52 *Reliable Distributed Systems*, ser. SRDS '07. Washington, DC, USA:
53 IEEE Computer Society, 2007, pp. 219–230. [Online]. Available:
54 <http://dl.acm.org/citation.cfm?id=1308172.1308237>
- 55 [26] D. Ionescu, "Microsoft bans up to one million users from
56 xbox live," *PC World*, Tech. Rep., 2009. [Online]. Available:
57 http://www.pcworld.com/article/182010/xbox_users_banned.html
- 58 [27] Sony, "Information on banned accounts and consoles," Sony
59 consumer electronics, Tech. Rep., accessed on may 2015.
60 [Online]. Available: [https://support.us.playstation.com/app/answers/](https://support.us.playstation.com/app/answers/detail/a_id/1260/~information-on-banned-accounts-and-consoles)
61 [detail/a_id/1260/~information-on-banned-accounts-and-consoles](https://support.us.playstation.com/app/answers/detail/a_id/1260/~information-on-banned-accounts-and-consoles)
- 62 [28] Brian, "Nintendo starting to ban pirates from on-
63 line services on 3ds," *Nintendo everything*, Tech.
64 Rep., 2015. [Online]. Available: [http://nintendoeverything.com/](http://nintendoeverything.com/nintendo-starting-to-ban-pirates-from-online-services-on-3ds)
65 [nintendo-starting-to-ban-pirates-from-online-services-on-3ds](http://nintendoeverything.com/nintendo-starting-to-ban-pirates-from-online-services-on-3ds)
- 66 [29] Wikipedia, "An analysis of proposed attacks against genuinity
67 tests," Tech. Rep., accessed on May 2015. [Online]. Available:
68 [http://en.wikipedia.org/wiki/Warden_\(software\)](http://en.wikipedia.org/wiki/Warden_(software))
- 69 [30] M. Kiperberg, A. Resh, and N. J. Zaidenberg, "Remote Attestation
70 of Software and Execution-Environment in Modern Machines," in
71 *CSCloud*, 2015.