

SEVerity: Code Injection Attacks against Encrypted Virtual Machines

Mathias Morbitzer*, Sergej Proskurin[†], Martin Radev*, Marko Dorfhuber[†] and Erick Quintanar Salas*

*Fraunhofer AISEC

{mathias.morbitzer,martin.radev,erick.quintanar.salas}@aisec.fraunhofer.de

[†]Technical University of Munich

proskurin@sec.in.tum.de, marko.dorfhuber@tum.de

Abstract—Modern enterprises increasingly take advantage of cloud infrastructures. Yet, outsourcing code and data into the cloud requires enterprises to trust cloud providers not to meddle with their data. To reduce the level of trust towards cloud providers, AMD has introduced Secure Encrypted Virtualization (SEV). By encrypting Virtual Machines (VMs), SEV aims to ensure data confidentiality, despite a compromised or curious Hypervisor. The SEV Encrypted State (SEV-ES) extension additionally protects the VM’s register state from unauthorized access. Yet, both extensions do not provide integrity of the VM’s memory, which has already been abused to leak the protected data or to alter the VM’s control-flow.

In this paper, we introduce the *SEVerity* attack; a missing puzzle piece in the series of attacks against the AMD SEV family. Specifically, we abuse the system’s lack of memory integrity protection to inject and execute arbitrary code within SEV-ES-protected VMs. Contrary to previous code execution attacks against the AMD SEV family, *SEVerity* neither relies on a specific CPU version nor on any code gadgets inside the VM. Instead, *SEVerity* abuses the fact that SEV-ES prohibits direct memory access into the encrypted memory. Specifically, *SEVerity* injects arbitrary code into the encrypted VM through I/O channels and uses the Hypervisor to locate and trigger the execution of the encrypted payload. This allows us to sidestep the protection mechanisms of SEV-ES. Overall, our results demonstrate a success rate of 100% and hence highlight that memory integrity protection is an obligation when encrypting VMs. Consequently, our work presents the final stroke in a series of attacks against AMD SEV and SEV-ES and renders the present implementation as incapable of protecting against a curious, vulnerable, or malicious Hypervisor.

I. INTRODUCTION

Cloud computing has become omnipresent and continues to get increasingly popular [1]. Despite its many advantages, in particular enterprises are reluctant towards entrusting data to provider-controlled infrastructures [2]. This lack of trust is twofold as the valuable data hinges on the integrity of the virtualization software, called the Hypervisor (HV), as well as the provider’s trustworthiness. In this regard, the HV’s integrity can be undermined by adversaries, despite the added isolation and security guarantees of the virtualization technology [3]–[6]. For instance, in the presence of vulnerabilities, skillful adversaries manage to escape their Virtual Machines (VMs) to take control over the HV and all VMs that share the same system [7]–[10]. Similarly, attackers with

physical access to the machine can run a series of attacks that reveal the VM’s memory contents [11]–[13]. Besides, in modern cloud infrastructures the cloud provider is able to expose potentially sensitive information [14]–[16]. As such, enterprises have to trust the cloud provider to preserve data integrity and confidentiality [17].

To counteract these concerns, AMD has announced Secure Encrypted Virtualization (SEV) [18]. SEV lends VMs the ability to encrypt their memory using a key maintained by a dedicated *Secure Processor (SP)*. Consequently, without knowing the respective secret key, neither the attacker performing cold boot attacks nor the HV will be able to reveal the VM’s memory contents. To also reduce information leakage through the VMs’s register state, AMD introduced *SEV Encrypted State (SEV-ES)* [30]. Unfortunately, it remains susceptible to attacks which do not rely on the VM’s registers. Another extension, *SEV Secure Nested Paging (SEV-SNP)*, aims to provide memory integrity protection. However, SEV-SNP is not yet available in hardware. Consequently, the missing memory integrity protection in SEV-ES continues to jeopardize sensitive data inside encrypted VMs.

In this paper, we present *SEVerity*, a new attack against SEV-ES. *SEVerity* allows a malicious HV to inject *truly arbitrary* code into the VM’s encrypted memory and cause the VM to execute the injected payload. Contrary to previous work, *SEVerity* does not limit itself to software bugs [27] or any memory encryption algorithm that is specific to a CPU version [23], [25], [26]. Further, *SEVerity* does not require any out-of-band collected knowledge of in-VM applications [20]–[25], [31]. In fact, our attack does not rely on any in-VM invariants, except for exported kernel symbols. Instead, it abuses the I/O channel to inject and execute truly arbitrary payloads in the VM’s encrypted memory.

To mount the *SEVerity* attack, first, we identify the VM’s page frame holding code that we can trigger from the outside of the VM, our *trigger page*. Next, we send a network packet with our payload to the VM. Since SEV-ES cannot directly pass packets into the encrypted memory, we use virtio buffers to track our payload in the VM’s memory. Finally, we abuse the lack of Second Level Address Translation (SLAT) integrity protection by remapping the trigger page to the page holding the injected payload, and trigger its execution.

In summary, we make the following contributions:

- We present SEVerity, the first attack against AMD SEV-ES which allows to inject and execute arbitrary code in an encrypted VM with kernel privileges, with no currently available countermeasures.
- We abuse SEV-ES' lack of SLAT integrity protection and its inability to grant DMA to encrypted memory to run SEVerity, without relying on flaws of the encryption algorithm's tweak function or firmware vulnerabilities.
- We evaluate SEVerity on all Linux distributions that are officially supported by SEV-ES and manage to achieve an attack success rate of 100 %.

Note that we have responsibly disclosed an early version of this paper to AMD. AMD has documented the attack vector under CVE-2020-12967.

II. BACKGROUND

Full disk encryption has gained popularity. Yet, even on-disk-encrypted data gets exposed when loaded into main memory [33]. As such, sensitive information, including passwords, cryptographic keys, or health related data can become target to attacks against the main memory [13], [34].

In this section, we outline AMD's hardware extensions for main memory encryption introduced in response to the depicted threat. We focus on AMD's efforts to encrypt the VM's memory and highlight the associated limitations. Particularly, we amplify the lack of Direct Memory Access (DMA) to the encrypted memory, which is key for SEVerity.

A. AMD Memory Encryption

AMD SME: To thwart attacks against the system's main memory, AMD has introduced Secure Memory Encryption (SME) [18], a hardware extension for both full and partial memory encryption. This extension links the system's page tables with its encryption engine, allowing to regulate which memory regions to encrypted. Specifically, AMD dedicates the 47th bit, the *C-bit*, of the system's Page Table Entries (PTEs) to determine whether the associated page should be encrypted. A dedicated ARM-based processor, the Secure Processor (SP), takes over the management of the secret key. The secret key used by the AES engine is generated on each system reset and solely known to the SP. By setting the *C-bit* of a particular PTE, the SP encrypts the respective page via a specially tweaked, high performance AES implementation [18]. The tweak incorporates the data's physical address into the encryption to prevent attackers from moving cipher-text blocks. Yet, SME has drawbacks when deployed in cloud environments. For one, the memory of VMs and the underlying HV is encrypted with the same key. Thus, curious administrators can access the sensitive data of VMs, despite full memory encryption. Additionally, having gained access to the encrypted memory of one VM, a local attacker will be able to similarly access the memory of all remaining VMs running on the same system.

AMD SEV: AMD Secure Encrypted Virtualization (SEV) draws upon SME [18]. Contrary to SME, SEV establishes additional isolation by assigning one secret key per isolation domain, including different VMs and the HV itself. Trying to access encrypted memory from a different domain associated with a different key returns garbled information. As such, the physical memory of isolation domains can only be decrypted with the secret key of the domain to which the memory has been assigned. By setting the *C-bit* of PTEs inside the VM's page tables, every VM can individually control which page to encrypt. If the VM sets the *C-bit* for a page, its contents will be encrypted with the VM's key. Only the associated VM can access the contents of such pages, which are referred to as *private* pages. By omitting the *C-bit*, the VM defines pages which can be shared with the HV. Depending on the *C-bit* in the HV's page tables, such *shared* pages can either be encrypted with the HV's key, or remain unencrypted. In both cases, the VM and the HV are able to access shared pages.

As the secret keys are governed by the SP and are inaccessible to the HV and VMs, there is no way to leak sensitive information maintained in different isolation domains. Yet, potentially sensitive data can still leak through register contents to the underlying HV [22].

AMD SEV-ES: In 2017, AMD has announced an iterative extension of the SEV, SEV Encrypted State (SEV-ES) [30], to prevent HVs from snooping or modifying the VM's register contents. SEV-ES additionally lends VMs the capability to encrypt and integrity-protect selected register state before handing over control to the HV, rendering the state inaccessible to the HV.

Even though SEV-ES overall manages to reduce the attack surface of VMs, its architecture neglects critical components prone to attacks. For instance, similar to SEV, SEV-ES does not protect the SLAT, which translates Guest Frame Numbers (GFNs) to System Frame Numbers (SFNs) [20], [24]. Also, SEV-ES denies DMA to the VM's memory, which must be established through alternative I/O channels [18]. It is precisely these channels which facilitate a severe attack vector that allows adversaries to inject and execute arbitrary code inside an encrypted VM, the analysis and disclosure of which is the main focus of our work.

Note that our work focuses on SEV-ES—the latest available version of the SEV family; SEVerity similarly applies to SEV, which faces the same fundamental security issues.

B. DMA into Encrypted Virtual Machines

Modern I/O Memory Management Unit (IOMMU) architectures increasingly focus on virtual environments. One task of IOMMUs is to transparently isolate direct accesses between devices and different VMs [35]. Unfortunately, this is not true for SEV-ES, as it prohibits DMA into VM-encrypted memory [18]. To bypass these limitations, SEV-ES makes use of a software IOMMU.

Generally, Linux implements an abstract DMA interface, hiding details of the low-level implementation. This grants Linux the flexibility to fall back to the software implementation of the IOMMU functionality, referred to as Software I/O Translation Buffer (SWIOTLB). SWIOTLB allows to emulate the DMA interface on systems which do not implement a physical IOMMU. Internally, SWIOTLB implements a *bounce buffering* mechanism to enable legacy devices to overcome the 4 GB boundary by granting the device access to the *bounce buffer*. Once the device has written to the bounce buffer, the system copies the contents to the buffer maintained by the device driver.

To overcome the limitations of SEV-ES regarding DMA, VMs use a similar approach. On read requests on non-protected VMs, the HV writes data directly into a buffer in the VM's memory from which the VM's device driver consumes it. Yet, SEV-ES prevents the HV from reading or writing meaningful data into the VM's memory. Instead, the HV copies the data into a buffer on a page shared between the VM and the HV. Then, the VM copies the data from the shared page to its destination in the VM's private memory. During this operation, the memory controller automatically encrypts the data when writing it to private memory.

Consequently, data bounced between the VM's shared and private memory is accessible by the HV. This leads to a window of opportunity for adversaries to inspect or modify the transferred information. Throughout this paper, we abuse the bounce buffering mechanism to inject arbitrary code into the VM's memory, without having to know the VM's encryption key. By combining such injections with SEV-ES' lack of memory integrity protection, we defeat AMD's efforts against code execution attacks.

III. ADVERSARIAL CAPABILITIES

To mount the SEV_{erity} attack, we assume that the attacker is in control of the HV. Whether the HV has been compromised or intentionally set up for malicious purposes is out of scope. The goal is to inject and execute arbitrary code in VMs, whose memory is encrypted by the latest AMD SEV-ES technology. This architecture generally forbids executing any code from shared, non-encrypted memory [36], preventing the injection and execution of unencrypted payloads. Further, the attacker is neither in the possession of the secret key nor has expert knowledge about the applied encryption algorithm. As such, the adversary cannot encrypt the payload before injecting it into the VM. The attacker is also not aware of any flaws in the utilized tweak function, which prevents moving cipher-text blocks in memory [23], [25], [26] (§ II-A). In other words, contrary to SEV_{erity} [26], the attacker is not concerned about the exact CPU version.

Even though the attacker does not require the source code of the kernel running inside the VM, we do assume access to the VM's kernel, or the kernel's exported symbol information. Inside the target VM, we expect a Linux

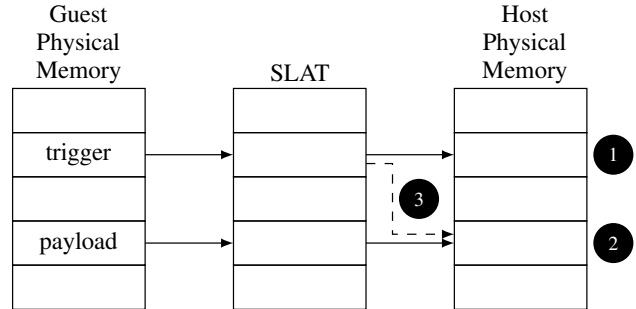


Figure 1: Overview of the SEV_{erity} attack. Having determined a trigger (❶), we send a payload to the VM (❷). After having determined the location of the payload, we modify the SLAT tables (❸) to trigger the execution of the payload.

kernel with Kernel Address Space Layout Randomization (KASLR). Also, we expect the VM to employ virtio for communicating with the system's devices. Both assumptions do not rely on any specific version of the Linux kernel or virtio. Further, disregarding the focus of the prototype presented in this work, the attacker is not limited to exploiting the network communication channel. In this context, the attacker disregards any in-VM IP or port filtering mechanisms, and is also agnostic to any in-VM services or open ports [20], [21], [23]–[25]. The attacker can inject truly arbitrary payloads into the VM's kernel, and therefore does not have to consider any kernel-level defenses that target attacks from user space. Despite knowing the addresses of the kernel's symbols, the attacker will have to determine the randomized kernel base address, imposed by KASLR [37]. Finally, the attacker relies on AMD's event injection interface to execute virtual interrupts at known registered locations.

IV. THE SEV_{erity} CODE EXECUTION ATTACK

The SEV_{erity} attack challenges the security model of AMD's state-of-the-art VM memory encryption feature SEV-ES [30]. For the attack, we expect the attacker to control the HV, recreating AMD's original attacker model, against which SEV-ES is intended to protect. This control allows the attacker to abuse architectural properties of SEV-ES to mount *arbitrary code execution attacks* against VMs in different memory isolation domains, despite the added confidentiality guarantees.

Our attack comprises the following three steps (Figure 1), which we elaborate in detail in this section:

- ❶ Identify a code region inside the VM, the execution of which we can trigger from the outside.
- ❷ Send a network packet with malicious payload to the VM and identify its location inside the VM's memory.
- ❸ Modify the Second Level Address Translation (SLAT) tables to point the identified trigger to the injected payload and trigger its execution.

To identify our trigger (❶), we remove the `execute` flag from all of the VM's physical memory pages, the Guest Frame Numbers (GFNs), and inject a Non-Maskable Interrupt (NMI) into the VM. This allows us to determine the GFNs that hold the NMI handler and the `do_nmi()` function, the latter of which we intend to use as trigger. To identify the location of network packets within the VM's memory (❷), we track the write accesses of the VM when copying the packet from the shared into private memory. Combined with detailed knowledge about `virtio`, this allows us to determine the exact location of every incoming network packet within the VM's memory. To conclude our attack, we inject a network packet containing our payload at a carefully calculated offset into the VM. Afterwards, we abuse the lack of integrity protection and modify the SLAT tables managing the mapping between GFNs and SFNs (❸). Specifically, we map the GFN containing the trigger to the SFN of the payload. Finally, by injecting an NMI into the VM, we trigger the execution of our payload.

Note that we do not limit SEVerity to `virtio` or Linux. Even though our work uses Linux and `virtio` as concrete examples to mount the SEVerity attack, the pillars upon which we base the attack are completely Operating System (OS) and subsystem (i.e., driver) independent. Thus, targeting another OS or subsystem would not change the essence of the three steps (❶ – ❸), which form the SEVerity attack (§ VII).

A. Identifying the Trigger Point

The virtualization extensions of AMD allow the system to pass interrupts directly to the VM. Alternatively, the HV can intercept an interrupt and subsequently inject a *virtual interrupt* into the VM through the *event injection* interface [36]. Just as real interrupts, virtual interrupts cause VMs to pause the currently executing code and handle the incoming event urgently in the associated interrupt handler. Consequently, virtual interrupts present themselves as good candidates that we can abuse for triggering code inside a VM. Yet, most interrupts are vital to the VM's execution and temporary replacing the interrupt's code with our payload would crash the VM. Therefore, we resort to one specific interrupt type, the NMI, that is rarely generated.

The CPU uses the Interrupt Descriptor Table, a hardware-defined data structure that registers all interrupt vectors, to locate the NMI handler. Once the HV injects a virtual NMI into a VM, the hardware immediately interrupts the VM's execution and invokes the NMI handler. In the context of SEVerity, we can abuse this behavior to initiate the execution of arbitrary code residing at the location of the registered NMI handler. Specifically, if we remap the GFN of the NMI handler to the SFN containing our payload (Figure 1) and issue a virtual NMI, we will be able to execute our payload.

Consequently, in order to perform the remapping, we have to identify the GFN of the NMI handler (❶). For this, we assume that the HV has access to the exported

symbol information or to the unencrypted image of the VM's kernel. This is a valid assumption, as VM images are often provided by the cloud provider and the unencrypted kernel image usually resides in `/boot`. By analyzing the kernel image, we can determine the offset of the NMI handler. We discuss alternatives for determining the offset in Section § VII. However, Linux applies KASLR to randomize the offset of the kernel image in the VM's virtual and physical address space [37], [38]. Fortunately, we are able to use different approaches that allow us to bypass KASLR and hence determine the GFN of the NMI handler. For instance, we can employ any of the following methods:

Manipulating KASLR: On modern AMD platforms, the Time Stamp Counter (TSC) counts the number of processor-clock cycles since the last reset [36], which can be read via the `RDTSC` instruction. The Linux kernel uses this instruction to create entropy at boot time, which is among others required to initialize KASLR. Therefore, a malicious HV manipulating the results of the `RDTSC` instruction at boot time will be able to pin the KASLR offset [27], allowing us to guess the KASLR offset used by the VM. However, this method requires a malicious HV to actively interfere with the VM at boot time. Additionally, the VM would be able to detect the static KASLR offset. To circumvent these limitations, we can also use techniques which allow us to infer the KASLR offset at runtime.

Fingerprinting nested page faults: Wilke et al. [26] observed that nested page faults occur *deterministically* during the boot process. Specifically, in the initial stage, before the VM's OS activates KASLR, the VM exhibits the same GFN access pattern across reboots. Only once the boot process starts to decompress the kernel image to a random offset in the VM's physical memory, the sequence will start to differ. We have determined that the first GFN which differs from this sequence holds the beginning of the kernel image in the VM's physical memory. As the beginning of the kernel image is always aligned on a 2MB boundary [39], we can safely assume that the start of the image will always remain at page offset zero. This method allows the HV to passively analyze the GFN access pattern of the VM's early boot stage. Yet, since the attacker does not necessarily have the chance to observe the bootstrapping phase of the VM, we have developed the following alternative method.

Probing the NMI handler: We can leverage the fact that we control the SLAT tables to identify the GFN of the NMI handler. Specifically, by revoking the `execute` permission from all GFNs in the SLAT tables, we can track all execution attempts of the VM. This configuration of the SLAT tables causes the VM to trap into the HV on every instruction fetch. For every trap, we can determine the GFN which caused the trap. That is, if we apply this configuration and inject a virtual NMI into the VM, the VM will attempt to execute the registered NMI handler and immediately trap

into the HV. Using this method, we are able to locate the GFN of the NMI handler. In order to continue the VM's execution, we restore the permissions of the GFN containing the NMI handler and resume the VM. Shortly after resuming the VM, it will trap again upon calling the `do_nmi()` function responsible for dispatching NMI's. Since `do_nmi()` resides in a different section (`.text`) than the NMI handler itself (`.entry.text`), we can be sure that the call to `do_nmi()` will trap into the HV. In other words, both functions reside on different GFNs and thus guarantee two consecutive traps.

Once we have identified the GFN holding the `do_nmi()` function, all that is left to determine is the function's exact offset within the respective GFN. This knowledge allows us to inject our payload at the same offset, ensuring it will be immediately executed after we modify the SLAT tables and trigger its execution. To determine the offset of the `do_nmi()` function, we analyze the exported symbol information of the VM's kernel image.

B. Identifying the Payload Destination

The architectural constraints of SEV-ES prohibit devices to apply DMA into encrypted VM memory (§ II-B). This limitation applies to physical as well as emulated devices, forcing VMs and devices to exchange data through shared pages. Even though this security-driven design decision aims to protect the system against DMA attacks [11], it also establishes a vantage point: the unprotected DMA channel facilitates attackers to (i) extract or inject arbitrary data from or to the VM, or (ii) inject arbitrary code into the VM's memory without knowing the secret encryption key (§ II-A). While previous attacks have targeted the former [23], the latter attack vector is the main focus of SEVerity. Specifically, we abuse the exposed network communication channel to inject arbitrary code into the encrypted memory. Yet, before injecting the payload, we need to identify where in memory the VM stores incoming network packets (⊙).

1) *Virtio tracking*: The goal of virtio tracking is to determine the GFNs in which the VM maintains incoming network packets. For this, we record the sequence of GFNs the VM is writing to when bouncing an incoming network packet from shared into private memory. To determine which of the GFNs contains which network packet, we need to understand how SEV-ES protected VMs use virtio [32] to exchange data between the virtual device in the HV and the driver in the VM. Virtio uses `virtqueues` as its main communication structures. Each virtual network device uses at least two `virtqueues`, one for receiving data (`rx virtqueue`), and one for sending data (`tx virtqueue`). The `virtqueues` can be packed or split. For simplicity reasons, we will focus on split `virtqueues`, which were used by all VMs we analyzed. Each split `virtqueue` comprises three components: (i) the *descriptor table*, (ii) the *available ring*, and (iii) the *used ring*.

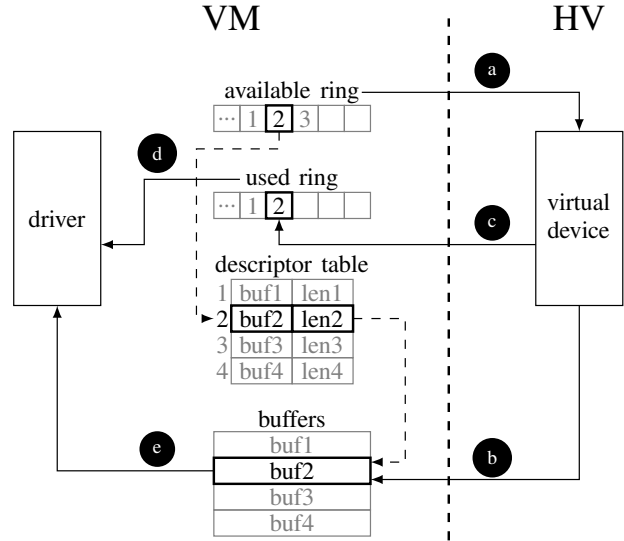


Figure 2: Abstract virtio network architecture. The virtual device in the HV leverages the *available ring buffer* and the *descriptor table* to locate the buffer for incoming packets. In the VM, the driver uses the *used ring buffer* and the *descriptor table* to determine the buffers holding the packets.

Each column of the descriptor table contains, among others, the address and length of a buffer the device and the driver can use for exchanging data. To refer to a buffer, the driver and the device use the buffer's position in the descriptor table as an index. The VM's driver writes indices into the available ring, offering buffers to the virtual device in the HV. Once the virtual device has used one or more of the available buffers, it adds the buffer's index to the used ring.

Figure 2 depicts the process of an incoming networking packet using the `rx virtqueue`. To enable the VM to receive network packets, the driver allocates buffers, adds them to the descriptor table and their respective index into the available ring. When a network packet comes in, the virtual device reads the available ring to find the next free buffer (a) and writes the packet into it (b). Then, the device adds the buffer's index to the used ring (c). By reading the used ring (d), the driver in the VM determines from which buffer to consume the network packet (e).

The described method requires the virtual device to access different parts of the `virtqueue`. As it is managed by the driver in the VM, the buffers are located within the VM's memory region. However, when using SEV-ES, the virtual device in the HV is not able to read or write meaningful data due to the encrypted memory. SEV-ES works around this limitation by using bounce buffers (§ II-B).

Using bounce buffers, the VM allocates the `virtqueue` structure in a memory region shared between the VM and the HV. After the device puts a packet into a shared buffer, the VM bounces the buffer's content from the shared,

unencrypted memory into the private, encrypted memory.

We track page accesses to determine to which location in private memory the VM bounces incoming network packets. Specifically, after the device in the HV places a network packet into a buffer in shared memory, we modify the SLAT tables by removing the `present` flag from the GFN holding the buffer. Hence, the HV gets notified as soon as the VM tries to access the buffer. Having received this notification, we conclude that the VM is about to read the buffer. Next, we enable *write tracking* on all other GFNs by revoking their write access permission in the SLAT tables. This allows us to collect the sequence of GFNs the VM is writing to until it has finished bouncing the packet into the private memory.

We abuse `virtio`'s notification system to determine when the VM has finished bouncing the packet. The HV raises an interrupt to notify the driver that it can consume a buffer. Yet, if the driver is busy, the `virtio` specification allows the driver in the VM to let the HV know that it does not want to receive interrupts. In older `virtio` implementations, the driver informs the device by setting the `VRING_AVAIL_F_NO_INTERRUPT` flag; the driver will unset this flag once it has finished processing the data. In recent `virtio` implementations, once the driver has finished consuming a buffer, it updates the `used_event` to the index of the consumed buffer. We monitor the `VRING_AVAIL_F_NO_INTERRUPT` flag and the `used_event` entry in order to comply with all `virtio` versions. This allows us to determine when the driver has completed processing the packet.

2) *Packet buffer tracking*: Having determined the write sequence when bouncing a network packet into the VM, we identify which GFN in the sequence contains the packet, and at which offset the packet is located. To determine the exact position of a packet, we have to understand how `virtio` manages its buffers. The buffers are allocated in 32KiB chunks of (physically) contiguous memory called a `packet buffer`. Each `packet buffer` is divided into multiple fragments of the same size. Figure 3 shows the layout of a `packet buffer` holding 21 fragments of 0x600 bytes each. The fragments are aligned and padded to 0x600 bytes, with the first fragment starting at offset 0x0. The dashed lines indicate the eight GFNs holding the buffer.

Note that within one `packet buffer`, all fragments have the same size. Yet, the fragment size can be dynamically adjusted to consider the size of recent network packets [40]. Therefore, the number of fragments per `packet buffer` can vary. On our VMs, we have observed a fragment size of 0x600 bytes. This size can change when the VM starts to receive larger frames. We can detect changes to the size of the fragments in the `packet buffer` by monitoring the length of the buffers in the descriptor table.

This knowledge of the `packet buffer`'s structure allows us to infer which GFNs contain which fragments, and which fragment starts at which offset. First, we know that the




Figure 3: The 32KiB sized `packet buffer` contains 21 fragments of 0x600 bytes. The size of 32KiB requires 8 GFNs, indicated by the dashed lines.

`packet buffer` is 32KiB aligned in the VM's memory. Therefore, we search the write sequences performed by the VM when bouncing a packet into private memory for a 32KiB aligned GFN. This GFN represents a candidate for the first GFN of a newly allocated `packet buffer`. Knowing that the first GFN contains the start of three fragments (Figure 3), we confirm the candidate by analyzing if the VM also writes to the GFN when bouncing the next two packets into private memory.

We can repeat this process for the subsequent GFNs in the `packet buffer` to minimize false positives. For this, we again use our detailed knowledge of the `packet buffer`. For example, the second GFN of the `packet buffer` contains the start of fragments 3, 4 and 5. Therefore, assuming we have correctly identified the first GFN of a `packet buffer`, the second, subsequent GFN should be accessed in the fourth, fifth, and sixth bouncing sequence. The third GFN containing the start of fragments 6 and 7 should be accessed in the sixth and seventh bouncing sequence, and so on.

Note that depending on the size of the incoming network packet, the VM may also access the second GFN in the third sequence. Considering the minimum size of an Ethernet frame (0x40 bytes) [41], the packet stored in the third fragment will not cause a write access to the second GFN. Yet, a packet bigger than 0x400 bytes will cause the VM to also write data to the second GFN; the fragment offset (0xc00) plus size (0x400) cross the first GFN's boundary.

By applying these conditions and adapting them to reduce the number of false positives, we are able to successfully identify the `packet buffer` within the VM's memory. Our tests show that requiring the first ten write sequences of the bounce buffer to fulfill our conditions eliminates the chance for false positives on all tested VMs. Having determined the position of the `packet buffer`, we can also determine the offset of each fragment by knowing its size. This allows us to predict where the VM will store prospective incoming network packets.

Note, the VM copies incoming network packets into the `packet buffer` before it applies any processing or filtering. Thus, the VM stores packets sent to closed or even firewalled ports in the `packet buffer`. This allows us to send our payload to closed or filtered ports and to eliminate the need for any active network service running in the VM.

V. EVALUATION

To show that SEVerity is effective and OS-agnostic, we created a Proof of Concept (PoC) and evaluated its performance on different VMs.

A. System Setup

Our setup comprises a Debian 11 host running a custom-built Linux kernel $v5.6$ including the unofficial SEV-ES patches [30]. The host equips an 8-core EPYC 3251 CPU with 64 GB of RAM and uses QEMU $v4.2.50$ and KVM to virtualize the VMs to be attacked. Further, we have extended the host's kernel with the SEVered [42] patches and the SEVerity implementation to provide all the functionality necessary for the attack, such as functionality required to identify the `packet buffer`.

We have evaluated SEVerity by targeting SEV-ES protected VMs running the following five different Linux distributions officially supported by AMD [43]: SUSE Linux Enterprise Server (SLES) 15; RedHat Enterprise Linux (RHEL) 8; Fedora 29; Ubuntu 18.04; and openSUSE-Tumbleweed. To enable SEV-ES, all VMs were running a custom built kernel in version $v5.7.0$ with patches for SEV-ES support [30].

B. Precision Evaluation

To evaluate SEVerity's precision, we ran it 1000 times against all five VMs. Before every round of the attack, we sent a random number between 0 and 64 packets with random data to avoid any bias from previous rounds. Afterwards, in every round of the attack, we identified the GFN of `do_nmi()` by injecting an NMI and performing execute tracking (§ IV-A). To identify the `packet buffer` in the VM's memory, we sent dummy packets and analyzed the page fault sequence for bouncing the packets into the VM's private memory (§ IV-B). Once a series of ten page fault sequences fulfilled our conditions, we deemed the `packet buffer` identified (§ IV-B). While a smaller number may have increased the speed of the evaluation, the higher number allowed us to eliminate false positives. Having identified the `packet buffer`, we sent additional dummy packets if the next free fragment was not located on a suitable offset to inject our payload (§ IV-C). The payload itself executed a hypercall with exit reason `0xff`, and afterwards issued a `RET` in order to return to the NMI handler. Having performed our preparations, we remapped the GFN of `do_nmi()` to the SFN of the payload and triggered its execution. By monitoring for hypercalls with exit reason `0xff`, we determined successful execution of the payload. To conclude the round, we connected to the VM via SSH to ensure that it continued to behave as expected.

To perform the evaluation itself, we filtered any external traffic to the VM to maintain full control over incoming network packets. In a real-world setting, incoming packets would have allowed us to determine the exact destination

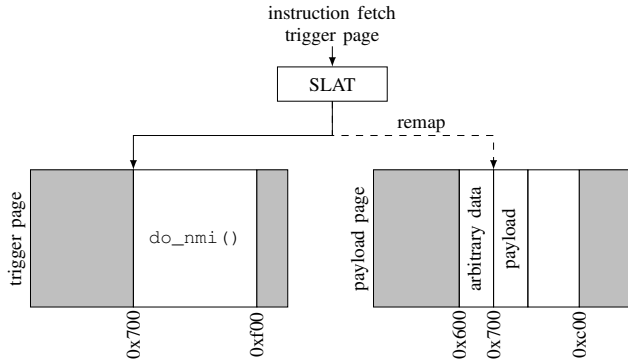


Figure 4: The system maps the GFN of the *trigger page* to the original SFN holding `do_nmi()`. SEVerity leverages the SLAT to remap the GFN of the trigger page to another SFN holding the injected payload. Thus, by injecting an NMI, we redirect the control-flow to our payload.

C. Injecting the Payload

Having determined the trigger point (§ IV-A) and the location to which to inject our payload (§ IV-B), we finalize the setup and conclude the SEVerity attack (⊕). The final step of the attack comprises two stages. First, we inject a payload into a suitable fragment in the `packet buffer`. Second, we manipulate the SLAT tables to remap the trigger such that its GFN translates to the SFN of the injected payload. Finally, we inject an NMI to execute the payload.

When injecting the payload into a suitable fragment, we have to consider that the trigger and the fragment can reside on different page offsets. As such, we must place the payload into a fragment whose location in memory spans a range comprising the needed page offset. In Figure 4, the trigger is located at the page offset `0x700`. Thus, we inject the payload to a fragment starting at offset `0x600`. To position the payload at precisely the same offset as the trigger, we fill the first `0x100` bytes of the network packet with arbitrary data. Since the fragment also holds protocol headers of the packet, the arbitrary data only spans the size of `0x100` bytes minus the size of the protocol headers (e.g., Ethernet frame, IP and TCP or UDP header). After a total of exactly `0x100` bytes, we place our payload to ensure that the payload and the trigger are located at the same page offset.

For the next step, we make use of the fact that while SEV-ES' memory encryption protects the guest page tables within the VM, the HV remains in charge of the SLAT tables. The SLAT tables map the VM's physical memory pages, the GFNs, to their actual location in the HV's physical memory, the SFNs. Therefore, once we prepared and injected the payload, we adjust the SLAT table entry such that the GFN that initially translated to the SFN holding the trigger (§ IV-A) translates to the SFN holding the injected payload (§ IV-B). That is, once the NMI handler attempts to call the `do_nmi()` function, it will instead execute our payload.

Table I: Evaluation details of 1000 runs of SEVerity against different VMs protected with SEV-ES. The table further summarizes an averaged number of the tracked write accesses to memory when bouncing network packets; the number of the sent packets; and the duration of the attack.

VM image	# Write Accesses	# Sent Packets	Duration	Success
SLES	38.52	19.97	2.57 sec	100 %
RHEL	66.78	19.50	2.75 sec	100 %
Fedora	68.74	20.37	3.04 sec	100 %
Ubuntu	43.10	19.89	2.65 sec	100 %
openSUSE	40.57	19.56	2.57 sec	100 %

of the payload in a stealthy way, without having to inject any artificial packets. Yet, to avoid having to rely on fluctuating numbers of incoming network packets, which would have influenced the evaluation results, we have opted for completely controlling the network traffic.

For each VM, we conducted our attack 1000 times, achieving a total success rate of 100%. Table I shows the detailed results of our evaluation. In the second column, we show the average number of write accesses recorded when monitoring the VMs bouncing a packet from shared to private memory. Their values range from 38.52 for openSUSE to 68.74 for Fedora.

In the third column, we show the average number of packets we had to send before we were able to identify the `packet buffer` in the VM's private memory. To ensure to always correctly identify the `packet buffer`, we required 10 consecutive write sequences to match our conditions. In our test environments, all VMs filled the `packet buffer` with 21 fragments of `0x600` bytes. Therefore, on average, we had to send 10 packets before a new `packet buffer` was allocated. The results in the third column match the expected value of these 10 packets plus 10 packets to ensure the `packet buffer` identification was correct.

The fourth column shows the average time to identify the `packet buffer` (⌚), inject our payload, remap and trigger its execution (⌚). The values are almost all below three seconds, ranging from 2.57s on SLES and openSUSE to 3.04s on Fedora. From these results, we conclude that while filtering external traffic during the attack may be noticeable, the short duration will not attract specific attention. Additionally, extending our implementation to buffer incoming packets for a short period of time instead of filtering external traffic would further increase stealthiness, as we would not need to send artificial packets. This leaves the VM with the possibility to analyze page access times to determine if the HV is performing memory access tracking. However, as we only require a maximum of around three seconds for our attack, it would be difficult to distinguish between delays caused by access tracking and other activities, such as migration of the VM.

Our results show that the identification of the trigger

page (§ IV-A) and the `packet buffer` (§ IV-B) work as described. This enables us to perform the SEVerity attack execute SEVerity against different VMs protected with SEV-ES with a success rate of 100%.

VI. COUNTERMEASURES

SEVerity mainly relies on the missing integrity protection to remap the GFN of our trigger to the SFN of our payload. Therefore, assuring the integrity of the VM's memory would prevent remapping, and be sufficient to prevent the SEVerity attack as well as others [20], [23]–[26], [44].

To ensure the integrity of the VM's memory, AMD has announced SEV-SNP [47] in 2020. One of the main goals of SEV-SNP is to provide integrity protection to the memory of SEV-protected VMs. SEV-SNP achieves this by adding a *Reverse Map Table (RMP)* to the address translation process. As with SEV-ES, the VM translates a virtual page to a GFN, and afterwards the HV uses the SLAT to translate the GFN into the SFN. In contrast to SEV-ES, with SEV-SNP, the SP afterwards verifies the integrity of the translation using the RMP. Each entry in the RMP contains among others information to which VM the SFN belongs to, the matching GFN and a *valid* flag. If an attacker modifies the SLAT tables to point a GFN to a different SFN, the SP will be able to detect this modification and the address translation will fault with a `#PF` exception. While SEV-SNP would therefore not prevent determining the location of our trigger and the `packet buffer` in memory, it would prevent mapping the GFN of the trigger to the SFN of the payload. Yet, up to this date, AMD has only published a whitepaper on SEV-SNP [47], without announcing when SEV-SNP will be released in hardware. At the same time, the currently available hardware, among others already used by cloud providers [48], is affected by SEVerity, posing a threat to users relying on the confidentiality and integrity properties of SEV-ES.

VII. DISCUSSION

Kernel image availability: In Section § IV-A, we assume access to the kernel image to extract the kernel's symbol information. This is a valid assumption as the kernel image typically resides in the unencrypted `/boot` partition of the VM's virtual disk image or, in some situations, can be directly provided by the HV. In both cases, the HV gains direct access to the kernel image. In case the VM stores its kernel image in an encrypted `/boot` partition, an attacker can use other methods to gain access to the VM's kernel image. For example, SEV-ES supports and disseminates only a small set of Linux distributions, each running a specific kernel version known to the public [43]. Hence, assuming the VM's provider did not compile a custom Linux kernel, the attacker can inspect the Linux kernel images to collect the necessary information to identify the trigger

point. Alternatively, to determine the used kernel version—or rather the addresses to relevant kernel functions—we could apply OS fingerprinting mechanisms [22], [31].

KASLR: We have proposed one approach to practically disable KASLR, and two approaches to identify the location `do_nmi()` within the VM’s memory with KASLR (§ IV-A). Both identification approaches work with the most recent implementation of KASLR. Yet, the granularity of KASLR might change in the near future. In a recent patch, Accardi [49] proposed to also randomize the offsets of functions within the kernel image. Using SEV-ES, after injection of an interrupt, the VM first internally handles the interrupt, which may cause it to access different GFNs before handing control to the NMI handler. To still be able to determine the location of the NMI handler and the `do_nmi()` page, we could analyze the page faults caused by the VM’s internal handler to determine when it finishes execution and hands control to the NMI handler. Alternatively, we could also resort to machine learning [31] to determine the location of the `do_nmi()` function.

Injected code size: Our PoC injects a small payload executing a hypercall that allows us to determine the success of the attack. Further, we also experimented with payloads up to 78 bytes. Yet, to inflict significant damage, the attacker would prefer to execute arbitrary code. Thus, the injected payload size could exceed the packet size limits imposed by the Maximum Transmission Unit (MTU). Yet, by injecting multiple payload fragments and using RIP-relative jumps, we can transfer the control-flow to other payload segments located in different fragments. This way, we can inject and execute payloads of a nearly arbitrary size.

Stealthiness: Our PoC is not designed to be stealthy. Yet, with some modifications, we can reduce its visibility. As part of SEVerity, we remap the GFN holding the `do_nmi()` function to the SFN holding our injected payload. Since we can cause the VM to execute the `do_nmi()` function from outside of the VM, it acts as a trigger for executing the injected payload. To still ensure consistent VM operation, we continue tracking the execution of incoming NMI requests by monitoring execute accesses to the GFN containing the NMI handler. That is, every time a regular NMI request occurs, we can restore the original mapping in the global SLAT tables and continue execution. Yet, walking the SLAT tables to change permissions of GFNs on every NMI event is a slow operation and can lead to race conditions in multi-virtual CPU (vCPU) environments. Besides, changes in the VM’s memory are visible to other vCPUs. To accommodate these issues, we can prepare a set of SLAT tables, which define different views (i.e., mappings) on the VM’s physical memory, and rapidly switch among them, without revealing the in-memory artifacts [6], [16], [50]. Since we can assign one SLAT table to each vCPU individually, the remaining vCPUs would not be able to detect the adjustments.

Further, even though our PoC blocks external traffic to the VM (§ V), we underline that this strategy is not a requirement. On the contrary, additional network traffic would allow us to detect the `packet buffer` in the encrypted memory. This way, we would not need artificial traffic to observe the sequences of write accesses of the bounce buffers. The same applies to identifying a suitable fragment in the `packet buffer`. Therefore, analyzing external traffic increases the stealthiness of the attack; essentially, the attacker would have to send only a single packet with the payload. To avoid issues caused by multiple network packets received shortly after each other, an improved implementation of our PoC can queue packets until analysis steps such as access page tracking of previous packets are finished. Therefore, SEVerity can also be applied to VMs with high network load. However, we opted for an evaluation without additional network traffic in our test environment, as this could have influenced our evaluation results.

Besides, in our PoC, we send UDP packets to a *closed* port. The VM can detect the unusual amount of traffic to a closed port and reveal the attack. To further increase the stealthiness of our approach, we could fall back to sending less suspicious packets, such as ARP packets. While ARP packets are limited in size, we can use them to detect the `packet buffer` and to fill unsuited `packet buffer` fragments with dummy packets.

Generalization: While we discussed SEVerity by injecting a payload via network packets, we neither rely on the usage of virtio or Linux, nor the network communication channel. Even though SEVerity uses Linux and virtio as concrete examples, the actual vulnerability is entirely independent from the OS, subsystem and driver. Assuming that another OS, such as Windows, would provide support for SEV-ES, the system would still be architecturally forced to transfer data from DMA-capable devices into the protected VM’s memory. The system will achieve this by copying the data from unencrypted to encrypted memory, which we exploit to inject our payload. Therefore, targeting another OS, subsystem or driver would not change the fundamentals of the three steps performed by SEVerity.

However, currently SEV-ES is only supported by Linux, and most Linux-based virtualization environments use virtio. For both Linux and virtio, our current implementation of SEVerity does not require a specific version. To be precise, by monitoring virtio’s `VRING_AVAIL_F_NO_INTERRUPT` flag as well as its `used_event` index, SEVerity can be applied to older as well as newer versions of virtio. Additionally, we tested SEVerity on different kernel versions between `v4.19.43` and `v5.7.0`. The kernel version is only limited by having to support SEV, which has been integrated into the Linux mainline kernel `v4.16` [43].

VIII. RELATED WORK

Hetzelt and Buhren performed the first in-depth analysis of SEV [24]. They made use of the lacking integrity protection and combined it with the HV's ability to modify the SLAT tables. By analyzing syscall sequences, the authors were able to modify the VM's login information for an SSH server, gaining unauthenticated access to the VM. In comparison to SEVerity, this attack requires interaction of the victim, which will happen only rarely on production systems. Another difference is that the attack requires an in-depth analysis of the target, for example by having access to a comparable VM. To perform such an analysis, the same authors also discussed to use machine learning to understand the behavior of the VM, which also requires a significant amount of preparation [31].

Morbitzer et al. [20] showed how a malicious HV can extract plaintext memory content from an encrypted VM. The attack uses a service in the VM to extract the content of arbitrary memory pages after modifying the SLAT tables. Compared to SEVerity, this approach requires a high number of requests. Therefore, Morbitzer et al. [21] improved the attack to perform targeted extraction of secrets from a VM. This drastically reduced the memory pages to be extracted, therefore also reducing the time required for the attack. A disadvantage that remains is that the target VM is required to provide an in-VM service, such as a web server.

Du et al. [25] analyzed the encryption algorithm of SEV. They determined that before encryption, the plaintext is tweaked with static values depending on the System Physical Address on which the data will be stored. This knowledge enabled them to gain access to an SSH server running within the encrypted VM by analyzing cipher blocks. The known tweak values were also used by Li et al. [23], who built a decryption oracle by replacing encrypted memory blocks about to be sent via an SSH connection with arbitrary blocks from the same VM. Using a similar technique, they were also able to create an encryption oracle. Building on the known tweak values, Wilke et al. [26] moved cipher blocks in the VM's encrypted memory, allowing them to execute code gadgets within the VM. However, all three attacks rely on knowing the tweak values to be able to meaningfully replace encrypted memory blocks in the VM's memory. While these tweak values were static on the first generation of EPYC CPUs, the second generation randomly generates the tweak values at every boot [26], making it resistant against the attacks. In comparison, SEVerity does not rely on knowing the tweak mechanism of the memory encryption to execute arbitrary code.

Werner et al. [22] monitored register values and memory accesses of an encrypted VM to derive and modify instructions the VM executed. Further, they used Instruction Based Sampling (IBS) to gather information about which applications were running within the encrypted VM. Compared to

their approach to read and modify confidential information via register values, their application fingerprinting technique can also be applied to SEV-ES. Therefore, it could be used as a base to implement OS fingerprinting to avoid the need to access the VM kernel's symbol information to determine the trigger function's page offset (§ IV-A).

Li et al. [44] presented an attack in which the HV modifies the Address Space Identifier of a VM during a `VMEXIT`, allowing them to decrypt the page tables of the target VM. By additionally modifying the VM's state, they created decryption and encryption oracles. As this modification of a VM's state is prevented by SEV-ES, their encryption and decryption oracles cannot be applied to VMs protected by SEV-ES. In comparison, SEVerity can be applied to SEV as well as to SEV-ES-protected VMs.

Radev and Morbitzer [27] manipulated the interfaces of a VM protected by SEV-ES. This allowed them to trick the VM into exposing confidential data to the HV, and to reduce the entropy of the VM's kernel probabilistic defenses. Additionally, they managed to execute code within the VM by providing false information about SEV support and adding additional MMIO regions. All of the proposed attacks make use of software bugs which have meanwhile been patched [45], [46]. In comparison, SEVerity uses a fundamental design flaw in SEV-ES, which will require hardware changes to be fully prevented.

IX. CONCLUSION

In this work, we have presented SEVerity, a powerful attack against AMD SEV-ES, which, compared to previous work, allows to inject and execute *truly arbitrary* code into encrypted VMs. The attack builds upon SEV-ES's lack of memory integrity protection and its inability to perform DMA into encrypted memory. Compared to previous work, SEVerity allows to inject and execute arbitrary code into protected VMs that cannot be prevented by currently available countermeasures. At the same time, we do not require any open ports in the VM, as we abuse the fact that Ethernet frames are pulled into the VM's memory before any filtering mechanisms can apply. Also, it suffices if our packets are copied into the VM's memory only for a short period of time.

We evaluated SEVerity against various VMs protected with SEV-ES. Our evaluation demonstrates a success rate of 100%, for which we need only a few seconds. Overall, our work is the next step in a series of attacks which highlight that AMD SEV-ES will not be able to provide sufficient security guarantees without protecting the integrity of the encrypted memory. Yet, we are optimistic that AMD's next evolution, SEV-SNP, will eliminate the weaknesses of its predecessors, providing stronger security guarantees to VMs in untrusted environments.

ACKNOWLEDGMENT

This work has been partially funded by the Fraunhofer Cluster of Excellence “Cognitive Internet Technologies”.

REFERENCES

- [1] A. DeNisco Rayome, “69% of enterprises moving business-critical applications to the cloud,” <https://www.techrepublic.com/article/69-of-enterprises-moving-business-critical-applications-to-the-cloud/>, January 2019, Accessed: 2021-27-01.
- [2] F. Amigorena, “Why SMBs Still do not Trust Cloud Storage Providers to Secure their Data,” <https://www.infosecurity-magazine.com/opinions/smb-trust-cloud-storage-1-1/>, 2019, Accessed: 2021-27-01.
- [3] Y. Liu, T. Zhou, K. Chen, H. Chen, and Y. Xia, “Thwarting Memory Disclosure with Efficient Hypervisor-Enforced Intra-Domain Isolation,” in *ACM Conference on Computer and Communications Security (CCS)*, 2015.
- [4] Q. Chen, A. M. Azab, G. Ganesh, and P. Ning, “PrivWatcher: Non-bypassable Monitoring and Protection of Process Credentials from Memory Corruption Attacks,” in *ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, 2017.
- [5] Z. Hua, D. Du, Y. Xia, H. Chen, and B. Zang, “EPTI: Efficient Defence against Meltdown Attack for Unpatched VMs,” in *USENIX Annual Technical Conference*, 2018.
- [6] S. Proskurin, M. Momeu, S. Ghavannia, V. P. Kemerlis, and M. Polychronakis, “xMP: Selective Memory Protection for Kernel and User Space,” in *IEEE Symposium on Security and Privacy (S&P)*, 2020.
- [7] VMware Inc., “VMSA-2017-0018.1,” <https://www.vmware.com/security/advisories/VMSA-2017-0018.html>, November 2017, Accessed: 2021-27-01.
- [8] SSD Secure Disclosure, “SSD Advisory – Oracle VirtualBox Multiple Guest to Host Escape Vulnerabilities,” <https://ssd-disclosure.com/ssd-advisory-oracle-virtualbox-multiple-guest-to-host-escape-vulnerabilities/>, January 2018, Accessed: 2021-27-01.
- [9] Citrix Systems, “Citrix Hypervisor Security Update,” <https://support.citrix.com/article/CTX263477>, 2019, Accessed: 2021-27-01.
- [10] —, “Citrix Hypervisor Security Update,” <https://support.citrix.com/article/CTX286756>, 2020, Accessed: 2021-27-01.
- [11] M. Becher, M. Dornseif, and C. N. Klein, “Firewire: all your memory are belong to us,” *Proceedings of CanSecWest*, 2005.
- [12] A. Boileau, “Hit by a bus: Physical access attacks with firewire,” *Presentation, Ruxcon*, vol. 3, 2006.
- [13] J. A. Halderman, S. D. Schoen, N. Heninger, W. Clarkson, W. Paul, J. A. Calandrino, A. J. Feldman, J. Appelbaum, and E. W. Felten, “Lest We Remember: Cold Boot Attacks on Encryption Keys,” in *USENIX Security Symposium*, 2008.
- [14] T. K. Lengyel, S. Maresca, B. D. Payne, G. D. Webster, S. Vogl, and A. Kiayias, “Scalability, Fidelity and Stealth in the DRAKVUF Dynamic Malware Analysis System,” in *Annual Computer Security Applications Conference (ACSAC)*, 2014.
- [15] S. Proskurin, J. Kirsch, and A. Zarras, “Follow the WhiteRabbit: Towards Consolidation of On-the-Fly Virtualization and Virtual Machine Introspection,” in *IFIP International Conference on ICT Systems Security and Privacy Protection (IFIP SEC)*, 2018.
- [16] S. Proskurin, T. Lengyel, M. Momeu, C. Eckert, and A. Zarras, “Hiding in the Shadows: Empowering ARM for Stealthy Virtual Machine Introspection,” in *Annual Computer Security Applications Conference (ACSAC)*, 2018.
- [17] F. Meixner and R. Buettner, “Trust as an integral part for success of cloud computing,” in *International Conference on Internet and Web Applications and Services (ICIW)*, 2012.
- [18] D. Kaplan, J. Powell, and T. Woller, “AMD Memory Encryption,” Advanced Micro Devices, Tech. Rep., 2016.
- [19] B. Singh, “x86: Secure Encrypted Virtualization (AMD),” <https://lwn.net/Articles/716165/>, March 2017, Accessed: 2021-27-01.
- [20] M. Morbitzer, M. Huber, J. Horsch, and S. Wessel, “SEVered: Subverting AMD’s Virtual Machine Encryption,” in *European Workshop on Systems Security (EuroSEC)*, 2018.
- [21] M. Morbitzer, M. Huber, and J. Horsch, “Extracting Secrets from Encrypted Virtual Machines,” in *ACM Conference on Data and Application Security and Privacy (CODASPY)*, 2019.
- [22] J. Werner, J. Mason, M. Antonakakis, M. Polychronakis, and F. Monrose, “The SEVerEST Of Them All: Inference Attacks Against Secure Virtual Enclaves,” in *ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, 2019.
- [23] M. Li, Y. Zhang, Z. Lin, and Y. Solihin, “Exploiting Unprotected I/O Operations in AMD’s Secure Encrypted Virtualization,” in *USENIX Security Symposium*, 2019.
- [24] F. Hetzelt and R. Buhren, “Security Analysis of Encrypted Virtual Machines,” in *International Conference on Virtual Execution Environments*, 2017.
- [25] Z.-H. Du, Z. Ying, Z. Ma, Y. Mai, P. Wang, J. Liu, and J. Fang, “Secure Encrypted Virtualization is Unsecure,” 2017. [Online]. Available: <https://arxiv.org/abs/1712.05090>
- [26] L. Wilke, J. Wichelmann, M. Morbitzer, and T. Eisenbarth, “SEVurity: No Security Without Integrity - Breaking Integrity-Free Memory Encryption with Minimal Assumptions,” in *IEEE Symposium on Security and Privacy (S&P)*, 2020.
- [27] M. Radev and M. Morbitzer, “Exploiting Interfaces of Secure Encrypted Virtual Machines,” in *Reversing and Offensive-oriented Trends Symposium*, 2020.

- [28] Cfir Cohen, “AMD-SEV: Platform DH key recovery via invalid curve attack: (CVE-2019-9836),” <https://seclists.org/fulldisclosure/2019/Jun/46>, June 2019, Accessed: 2021-27-01.
- [29] R. Bühren, C. Werling, and J.-P. Seifert, “Insecure Until Proven Updated: Analyzing AMD SEV’s Remote Attestation,” in *ACM Conference on Computer and Communications Security (CCS)*, 2019.
- [30] David Kaplan, “Protecting VM Register State with SEV-ES,” White Paper, 2017.
- [31] R. Bühren, F. Hetzelt, and N. Pirnay, “On the detectability of control flow using memory access patterns,” in *Workshop on System Software for Trusted Execution (SysTEX)*, 2018.
- [32] Edited by Michael S. Tsirkin and Cornelia Huck, *Virtual I/O Device (VIRTIO) Version 1.1*, OASIS Committee Specification 01, 2019, latest version: <https://docs.oasis-open.org/virtio/virtio/v1.1/virtio-v1.1.html>.
- [33] Y. Tang, P. Ames, S. Bhamidipati, A. Bijlani, R. Geambasu, and N. Sarda, “CleanOS: Limiting Mobile Data Exposure with Idle Eviction,” in *USENIX Symposium on Operating System Design and Implementation (OSDI)*, 2012.
- [34] R.-P. Weinmann, “Baseband Attacks: Remote Exploitation of Memory Corruptions in Cellular Protocol Stacks,” in *USENIX Workshop on Offensive Technologies (WOOT)*, 2012.
- [35] M. Ben-Yehuda, J. Mason, J. Xenidis, O. Krieger, L. Van Doorn, J. Nakajima, A. Mallick, and E. Wahlig, “Utilizing IOMMUs for Virtualization in Linux and Xen,” in *Ottawa Linux Symposium*, 2006.
- [36] Advanced Micro Devices, “AMD64 Architecture Programmer’s Manual (Volumes 1-5),” 2020.
- [37] J. Edge, “Kernel Address Space Layout Randomization,” <https://lwn.net/Articles/569635/>, 2013, Accessed: 2021-27-01.
- [38] The kernel development community, “Kernel Self-Protection,” https://www.kernel.org/doc/html/latest/_sources/security/self-protection.rst.txt, 2020, Accessed: 2021-27-01.
- [39] M. Schwarz, C. Canella, L. Giner, and D. Gruss, “Store-to-Leak Forwarding: Leaking Data on Meltdown-resistant CPUs,” 2019. [Online]. Available: <https://arxiv.org/abs/1905.05725>
- [40] M. Dalton, “[PATCH net-next 4/4] virtio-net: auto-tune mergeable rx buffer size for improved performance,” <https://lists.linuxfoundation.org/pipermail/virtualization/2013-November/025626.html>, 2013, Accessed: 2021-27-01.
- [41] “IEEE Standard for Ethernet,” *IEEE Std 802.3-2018 (Revision of IEEE Std 802.3-2015)*, pp. 1–5600, 2018.
- [42] M. Morbitzer and M. Huber, “Github - SEVered Framework,” <https://github.com/Fraunhofer-AISEC/severed-framework/>, 2019.
- [43] Advanced Micro Devices, “GitHub - AMDESE/AMDSEV: AMD Secure Encrypted Virtualization,” <https://github.com/AMDESE/AMDSEV>, 2018, Accessed: 2021-27-01.
- [44] M. Li, Y. Zhang, and Z. Lin, “CROSSLINE: Breaking “Security-by-Crash” based Memory Isolation in AMD SEV,” 2020. [Online]. Available: <https://arxiv.org/abs/2008.00146>
- [45] M. Radev, “[PATCH v7 72/72] x86/sev-es: Check required CPU features for SEV-ES,” <https://lkml.org/lkml/2020/9/7/1074>, 2020, Accessed: 2021-27-01.
- [46] J. Roedel, “[PATCH v2 0/5] x86/sev-es: Mitigate some HV attack vectors,” <https://lkml.org/lkml/2020/10/20/465>, 2020, Accessed: 2021-27-01.
- [47] Advanced Micro Devices, “AMD SEV-SNP: Strengthening VM Isolation with Integrity Protection and More,” <https://www.amd.com/system/files/TechDocs/SEV-SNP-strengthening-vm-isolation-with-integrity-protection-and-more.pdf>, 2020.
- [48] Google, “Confidential vm and compute engine,” 2021, <https://cloud.google.com/compute/confidential-vm/docs/about-cvm>.
- [49] J. Edge, “Finer-grained kernel address-space layout randomization,” <https://lwn.net/Articles/812438/>, 2020, Accessed: 2021-27-01.
- [50] Xen-devel mailing list, “Alternate p2m design specification,” <https://lists.xenproject.org/archives/html/xen-devel/2015-06/msg01319.html>, 2015, Accessed: 2021-27-01.