Extracting Secrets from Encrypted Virtual Machines

Mathias Morbitzer*
Fraunhofer AISEC
Garching near Munich, Germany
morbitzer@aisec.fraunhofer.de

Manuel Huber*
Fraunhofer AISEC
Garching near Munich, Germany
manuel.huber@aisec.fraunhofer.de

Julian Horsch
Fraunhofer AISEC
Garching near Munich, Germany
julian.horsch@aisec.fraunhofer.de

ABSTRACT

AMD SEV is a hardware extension for main memory encryption on multi-tenant systems. SEV uses an on-chip coprocessor, the AMD Secure Processor, to transparently encrypt virtual machine memory with individual, ephemeral keys never leaving the coprocessor. The goal is to protect the confidentiality of the tenants' memory from a malicious or compromised hypervisor and from memory attacks, for instance via cold boot or DMA. The SEVered attack has shown that it is nevertheless possible for a hypervisor to extract memory in plaintext from SEV-encrypted virtual machines without access to their encryption keys. However, the encryption impedes traditional virtual machine introspection techniques from locating secrets in memory prior to extraction. This can require the extraction of large amounts of memory to retrieve specific secrets and thus result in a time-consuming, obvious attack. We present an approach that allows a malicious hypervisor quick identification and theft of secrets, such as TLS, SSH or FDE keys, from encrypted virtual machines on current SEV hardware. We first observe activities of a virtual machine from within the hypervisor in order to infer the memory regions most likely to contain the secrets. Then, we systematically extract those memory regions and analyze their contents on-the-fly. This allows for the efficient retrieval of targeted secrets, strongly increasing the chances of a fast, robust and stealthy theft.

CCS CONCEPTS

• Security and privacy → Virtualization and security.

KEYWORDS

AMD SEV; virtual machine encryption; virtual machine introspection; memory extraction; data confidentiality

*Both authors contributed equally to the work.

1 INTRODUCTION

On common multi-tenant systems, the confidentiality of sensitive Virtual Machine (VM) data depends on both the Hypervisor's (HV) integrity and on the operator's trustworthiness. Unfortunately, these strong requirements are prone to getting infringed by different attack vectors. Examples are attacks by other tenants exploiting software-level vulnerabilities to escape their sandboxed VMs [15, 21, 23], attackers with physical access conducting a memory attack, e.g., via Direct Memory Access (DMA) [3, 4, 7] or cold boot [10], or simply a malicious operator using the HV to read the VM's memory. In order to protect the VM's memory in such scenarios, AMD introduced Secure Encrypted Virtualization (SEV) [2, 6] on recent server systems. SEV is a hardware extension for main memory encryption on a per-VM granularity. With SEV enabled, AMD's Secure Processor (SP) transparently encrypts the main memory of each VM with individual SP-bound keys. The goal is to protect VMs from memory attacks and from a malicious or compromised HV. To attest tenants that their VMs' memory is indeed encrypted, SEV includes a cryptographic protocol to verify VM encryption on an SEV-enabled platform.

SEVered [16] is a recent attack on AMD SEV, which showed that it is nevertheless possible for a HV to extract plaintext contents from SEV-encrypted VMs. SEVered exploits SEV's missing integrity protection for VM memory pages, previously discovered by [11, 17], to modify the memory mapping of a non-colluding service inside a VM. The modification causes the service to access and return an arbitrary portion of plaintext memory when serving requests. This allows an attacker in the HV to extract all the VM's main memory in plaintext by repeatedly requesting the same resource and changing its mapping. However, the encryption prevents the attacker from locating the VM's most valuable resources in memory prior to extraction, such as keys for Transport Layer Security (TLS), Secure Shell (SSH) or Full Disk Encryption (FDE). In the worst case, extracting those secrets requires a full dump of the VM's memory. This can take a significant amount of time, depending on the size of the attacker-controlled resource and throughput of the service. For example, an extraction speed of about 80 KB/s was reached with web servers providing a resource covering exactly one memory page. In this scenario, it takes more than 7 hours and requires 524,288 requests to extract all memory contents of a VM with 2 GB of main memory. During this time, other clients requesting the same resource also receive arbitrary contents, making full memory extraction conspicuous.

In this paper, we show that it is possible to overcome these downsides and present an approach that makes HVs capable of quickly locating and extracting specific secrets from SEV-enabled VMs. Our approach has two phases, the *observation* and the *retrieval* phase. In the observation phase, we exploit the fact that the HV is able to observe certain events triggered by VMs. These *observable*

events can, for instance, be page faults which the HV handles but also I/O events like network traffic or disk writes. We observe and combine such events to identify a minimal set of VM memory pages likely to contain the targeted secrets. Second, in the retrieval phase, we iteratively extract and analyze the identified set of pages on the fly until we find the targeted secret. For this phase, we use the SEVered attack, but could potentially leverage other vectors allowing memory extraction from SEV-encrypted VMs.

Our targeted extraction approach offers an inconspicuous, reliable and efficient method to steal various secrets from encrypted VMs. We demonstrate the potential of our approach by extracting TLS and SSH keys from a VM's user space memory, and FDE keys stored in the VM's kernel space. We conduct our experiments on an SEV-enabled EPYC processor, running Apache and nginx web servers as well as the OpenSSH server. To show that our approach can cope with real-world scenarios where VMs can be under varying levels of load, we base our experiments on a load model in which multiple independent clients concurrently access the VM's services during our attack.

2 AMD SEV AND THE SEVERED ATTACK

This section provides background information on AMD SEV, the Second Level Address Translation (SLAT) concepts of HVs, and the SEVered attack.

SEV. The AMD SEV technology allows for the transparent encryption of main memory of individual VMs. SEV primarily targets server systems and builds on the AMD Secure Memory Encryption (SME) technology, which provides transparent full main memory encryption. While the goal of SME is to protect systems against physical attacks on the main memory, SEV tries to additionally protect memory of individual VMs against attacks from other VMs and from a malicious HV. The SEV encryption is executed by a hardware AES engine located in the memory controller. The keys for the encryption are created and managed by an additional component, the AMD SP. All keys are ephemeral and never exposed to software on the main CPU. In contrast to SME, SEV uses different keys for each VM and for the HV. Additionally, a VM running on an SEV-protected system can request encryption and receive proof that its memory contents are being encrypted, which establishes trust between its owner and the remote operator. While SME was first integrated into AMD's Ryzen CPUs, SEV was introduced onto the market with the EPYC processor family. The mainline Linux kernel provides necessary software-level support for SEV.

SLAT. AMD SEV integrates with the existing AMD hardware virtualization technologies marketed as AMD-V. An integral component of hardware virtualization is an additional address translation, often named *nested paging* or SLAT [1]. While non-virtualized systems simply translate virtual addresses directly to physical addresses, a hardware-virtualized system distinguishes between three different types of addresses. When the VM accesses a Guest Virtual Address (GVA), the guest-controlled first level translation translates the address to a Guest Physical Address (GPA). The GPA is then translated to a Host Physical Address (HPA) using the second-level translation controlled by the HV. SLAT is completely transparent to the VM. This allows running multiple VMs that use the same GPA

space while separating them in physical memory. With SEV enabled, the first level translation from GVA to GPA in the encrypted VM is non-accessible to the HV. But the HV is still responsible for managing physical memory for its VMs and is therefore able to restrict access and change second-level mappings from GPAs to HPAs. Since there is no integrity protection in SEV, the HV can use SLAT to transparently switch a GPA to HPA mapping to a different HPA page belonging to the same VM.

SEVered. The SEVered attack [16] enables a malicious or compromised HV to extract the full memory of SEV-encrypted VMs in plaintext by exploiting SEV's missing integrity protection. SEVered requires a (non-colluding) service in the targeted VM, e.g., a web server, offering a remotely accessible resource. The first step of SEVered is to identify the HPAs, i.e., the physical pages, at which the accessible resource is located in the VM's encrypted memory. The number of pages containing (parts of) the resource depends on the size of the resource as well as on the page size. The knowledge about the resource's location allows SEVered to modify the VM's GPA to HPA mappings to point to arbitrary other HPAs of the VM instead of to the service's resource. The modified mapping causes the service to access different memory pages instead of the real resource when handling requests. In the second step, SEVered repeatedly requests the resource while remapping the memory (using the SLAT feature). This leads to the iterative extraction of an encrypted VM's full memory contents. The throughput of SEVered depends on the service and on the amount of pages that can be extracted at once. Our attack uses SEVered for the extraction of main memory from SEV-encrypted VMs. Like SEVered, our attack neither requires breaking SEV's cryptographic primitives, nor control over the SP. Likewise, our attack requires control over the HV, i.e., a malicious administrator or a compromise of the HV. We refer to [16] for further information about SEVered.

3 FINDING AND EXTRACTING SECRETS

Our concept for the targeted extraction of secrets from SEV-encrypted VMs has two phases: In the first phase, we start our attack by *observing* the page accesses of the targeted VM in the HV until an event occurs which indicates the VM's recent use of the targeted secret. In the second phase, we *search* the VM's memory for the secret by systematically extracting and analyzing the set of observed pages accesses. This section describes both phases in detail.

3.1 Observation Phase

The goal of the observation phase is to narrow down the set of VM memory pages possibly containing the targeted secret. We start the phase at an arbitrary point in time by tracking the VM's page accesses in the HV until observing the end of a particular *activity*. This activity *must* make use of the targeted secret *at least once*. The start of the activity, denoted by $Activity_{Start}$, does not need to be observable by the HV. In contrast, the end of an activity, called $Activity_{End}$, *must* be a HV-observable event. This event indicates that the VM *recently* used the secret one or multiple times, denoted by $Use_1..Use_n$. As soon as we observe $Activity_{End}$, we stop tracking, denoted by $Tracking_{End}$. We do not actively attempt to trigger $Activity_{Start}$ in order to interfere as little as possible.

To start page access tracking, denoted by $Tracking_{Start}$, the HV invalidates all the target VM's GPA to HPA mappings. As a consequence, each of the VM's page accesses causes an observable event, a SLAT page fault. For each SLAT page fault, we record the GPA as well as the time and type of the page access (read, write, execute) in a list and re-validate the mapping. The re-validation *clears* the page from tracking. This means that each accessed page triggers exactly one page fault and that we track the page *exactly once*, namely the first time it is accessed after $Tracking_{Start}$. The tracking enforces that accesses to the secret will inevitably be recorded. Note that the secret can be contained in a single page or span over multiple pages and can have multiple occurrences on different pages.

An example for an activity is a TLS handshake as part of a request to a web server. The server uses the targeted secret, in this case its TLS private key, to authenticate itself to a client during the handshake. The HV can observe $Activity_{End}$ by monitoring network traffic, waiting for the packet the VM sends to complete the handshake.

Figure 1 depicts an attack scenario with the target VM and the HV in the upper and lower box, respectively. The illustration shows the start and end of a VM's activity along with events triggered by the activity, such as $Use_1..Use_n$. The vertical arrows crossing the upper and lower box represent the events observable from the HV. These are, for instance, SLAT page faults, network packets or disk I/O. Some of the vertical arrows do not cross the boundary of the VM. These are events not observable by the HV, for example, page faults handled by the VM or possibly $Activity_{Start}$. Some of the events may be related to concurrent activities, and multiple other activities may potentially make use of the secret as well, cases which are not depicted in Figure 1. The illustration emphasizes that $Tracking_{End}$ concludes the observation phase right after $Activity_{End}$.

When starting tracking between Use_n and $Activity_{End}$, we do not observe any of the page accesses to the secret. This means that we are unable to find the secret in the later search phase, requiring to repeat the attack. This is why we call the timespan between Use_n and $Activity_{End}$ the *critical window*. The critical window size is an important factor regarding the quality of the attack. The smaller the critical window the higher the probability that the attack succeeds. Further, a small critical window means quick termination of page access tracking after Use_n. This causes Use_n to be tracked *at the very end* of the phase, and likely only a few more pages to be tracked after Use_n. We evaluate the critical window size for different scenarios with various levels of load in Section 5.

It is *not* necessary to synchronize the start of the observation phase with a possibly non-observable $Activity_{Start}$. If $Tracking_{Start}$ takes place *long before* $Activity_{Start}$, the observation phase might take longer, but since every page is tracked only once, this does not lead to a persistent performance impact. On the other hand, if $Tracking_{Start}$ takes place *after* $Activity_{Start}$ (but not inside the critical window), the tracking period will be shorter and likely output less tracked page accesses.

To conclude, the result of the observation phase is a list of pages in which the page with the targeted secret is contained at least once as long as $Tracking_{Start}$ is not inside the critical window. The set of pages in the list is significantly smaller than the whole set of the VM's pages.
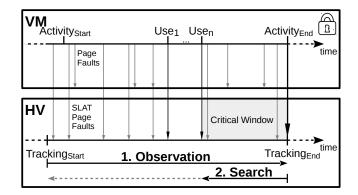

Figure 1: A HV first observing an activity inside an encrypted VM and then searching for the targeted secret. The vertical lines crossing the VM boundary into the HV box depict the events observable outside the VM.

3.2 Search Phase

The goal of the search phase is to extract the targeted secret from the VM's memory as quickly as possible, i.e., with a minimal number of memory requests. The input to the search phase is the list of tracked pages acquired during the observation phase. It is unknown which of the page accesses in the list correspond to $Use_1..Use_n$. The naive extraction of all pages in the list would still require a fairly high number of memory requests to find the secret. In the following, we describe our approach for a more efficient extraction.

The search phase starts right after $Tracking_{End}$, as depicted at the bottom of Figure 1. We know that $Activity_{End}$ indicates *recent* use of the secret. This means that Use_n must have occurred shortly before $Tracking_{End}$. For this reason, we consecutively extract the tracked pages *in backward order* until we find the secret. We thus start the extraction with the *most recently* tracked pages. This backward search is shown by the arrow directed to the left at the bottom of Figure 1. We analyze extracted memory chunks for the presence of the secret *on the fly* to be able to terminate the extraction procedure as early as possible. On the fly means we search the latest extracted memory chunk for the secret while we request the next chunk. When finding the secret in the chunk, we terminate the search phase, otherwise we request another chunk. The actual analysis is specific to the targeted secret and described in Section 4 for different secrets.

We propose an optional *preprocessing* step before the extraction to further minimize the number of memory requests. Preprocessing *filters* page accesses from the list, which *cannot* represent a use of the secret, and *prioritizes* accesses that are *likely* to represent a use. The ability to filter and prioritize depends on the use case, in particular, on the specific activity and secret. In most cases, the secret is a data structure on a page in non-executable memory, allowing to filter all execute-accesses from our list. The page is likely to be read, but may also be written during an activity. Depending on the use case, it is also possible that the secret resides in a read-only area, or represents confidential code. The information about this can often be acquired prior to the attack. A further possibility for preprocessing is to conduct a representative offline access pattern

analysis for the activity to observe the expected timing of $Use_1..Use_n$. An offline analysis is more representative the more the hardware platform and the software configuration inside the VM resemble the attack target. With the gained timing information, an attacker can further filter or re-prioritize pages in the list.

Extracting the secret from the encrypted VM using SEVered requires the secret to remain at the same location during the attack. This means that the secret must not be erased or moved to a different HPA by the VM's kernel before the search phase terminates. We show that the secrets we chose for extraction always fulfilled this requirement and investigate preprocessing possibilities as part of our evaluation in Section 5.

4 KEY EXTRACTION SCENARIOS

In the following, we describe the application of our concept for the extraction of targeted secrets at the example of private keys and symmetric FDE keys. We focus on the aspects from Section 3 that are specific for the type of secret. These aspects are the activities with their events and the on the fly analysis.

4.1 Private Keys

For the extraction of private keys, we focus on the example of web server TLS keys. These keys are resources located in a VM's user space and highly sensitive. Web servers use these keys to establish authenticated TLS channels with clients. An attacker can make use of a stolen private key for identity spoofing and deceive clients for fraud or data exfiltration.

Events. $Activity_{Start}$ is the start of a TLS handshake. The handshake can be part of an HTTPS request or be directly triggered by a client. Use_i represents a server's use of the TLS key for its authentication during the handshake. The exact moment of use depends on the key exchange method. For instance, in case of an Elliptic-Curve Diffie-Hellman Ephemeral (ECDHE)-based key exchange algorithm, this is the moment of signing curve parameters. For an RSA-based key exchange, this moment is the decryption of the premaster secret encrypted by the client with the server's public key. $Activity_{End}$ happens when the VM sends the client a specific network packet during the handshake. We observe these packets with network monitoring tools. The *change cipher spec* packet is an indicator independent of the specific key exchange algorithm. Depending on the algorithm, packets sent earlier may be usable indicators as well. Note that we can also observe or even trigger $Activity_{Start}$ ourselves in this scenario. We discuss this aspect in Section 6.

On the fly analysis. The public key and its length are part of the server's certificate and known in advance. When using RSA, the private components of the key are the factors p or q of known length dividing the modulus of the public key. For every extraction request we make, we traverse the extracted chunk of memory and check if it contains a contiguous bit sequence that divides the modulus without remainder. If so, we found either p or q and can instantly determine the other factor. Otherwise, we request the next chunk of memory. Analyzing a chunk this way usually takes less time than memory extraction with SEVered, see Section 5.

The same approach can be used for extracting SSH private keys. In the SSH scenario, the SSH server must also use its private key for authentication during the SSH handshake when establishing a session. We evaluate the extraction of TLS and SSH keys using the Apache, nginx and OpenSSH servers in Section 5.

4.2 FDE Keys

The normal approach when using SEV is to first perform an attestation of the platform. The attestation proves to the tenant that the VM has been started with SEV enabled. After a successful attestation, the tenant provides the FDE key in encrypted form to the VM [2]. This protects the key from eavesdropping adversaries in the network and from being read by the HV. Thereafter, the FDE key is present in the VM's memory and can be extracted with our approach. The FDE key is particularly important, because it allows attackers to decrypt the VM's persistent storage gaining access to further valuable secrets.

Events. The corresponding activity is a disk I/O operation. The trigger for $Activity_{Start}$ is not observable by the HV and unlike in the TLS key scenario, $Activity_{Start}$ can have many different triggers. The trigger can, for instance, be data uploaded to a service, a request to a web server being logged, or an operation of the VM's OS involving disk I/O. The event Use_i is the VM's use of the FDE key to en- or decrypt disk content to be read or written. We observe $Activity_{End}$ by monitoring the VM's disk image file in the HV.

On the fly analysis. We can be sure that we found the secret as soon as we are able to successfully decrypt the VM's persistent storage. Traversing extracted memory chunks and naively trying each possible sequence as key leads to an inefficient approach. Our goal is thus to first identify key candidates in extracted memory chunks. For this purpose, we search the extracted memory for characteristics specific to FDE keys based on the following two criteria.

First, the FDE key is stored in the VM's kernel in a specific data structure. This structure has various fields, some of which must have certain value ranges, for instance, kernel addresses pointing to other kernel objects. Our first criterion for a key candidate is thus the identification of possible FDE key structures in extracted memory chunks.

Our second criterion is based on the statistical properties of the FDE key. Because FDE is usually AES-based, the kernel derives round keys from the FDE key and keeps them in *AES key schedules* in memory. The round keys have common statistical properties that can be identified with linear complexity. The first-round key is the AES key itself. We use aeskeyfind [5] to search memory chunks for AES key schedules. Note that candidates that turn out to be false positives are possibly symmetric keys used for other purposes and might also be valuable secrets. The traversal of memory chunks based on these two criteria takes considerably less time than the extraction of memory with SEVered, see Section 5.

We evaluate the FDE key extraction scenario as part of the following section.

5 IMPLEMENTATION AND EVALUATION

In the following, we first define performance indicators and then present our prototype and test setup. Based on that, we evaluate the extraction of TLS, FDE and SSH keys, as discussed in Section 4. In the final part of our evaluation, we present strategies for optimization with preprocessing and summarize our results.

5.1 Performance Indicators

The key factors we investigate are the *success probability* and the *attack time*.

Success Probability. As discussed in Section 3, the critical window size is the factor determining the success probability of our attack. The smaller the critical window, the smaller the probability that the observation phase ends without having tracked the access to the secret. In our evaluation, we present the success probability for the tested scenarios and provide an upper bound on the size of the critical window. We call the upper bound the *reaction time* of our attack. The reaction time is the sum of the critical window (the time frame between Use_n and $Activity_{End}$) and the time our prototype requires to detect $Activity_{End}$ and stop tracking (the time frame between $Activity_{End}$ and $Tracking_{End}$).

Attack Time. We divide the total attack time of a full attack into its three components: the time required to setup SEVered prior to extraction, the duration of the observation phase and the duration of the search phase.

Setup of SEVered. The time required to setup SEVered is evaluated in [16] and is thus not subject of our evaluation. Setting up SEVered usually takes less than 20 seconds, depending on the load of the VM. After setting up SEVered once, we can arbitrarily extract the victim VM's memory and repeat our attack when necessary.

Observation phase. The main factor for the duration of the observation phase is the frequency of the targeted activity. For instance, a web server under high load will often make TLS handshakes while SSH logins generally occur less frequently.

Search phase. The duration of the search phase is mainly determined by the amount of memory that has to be extracted until the secret is found. This is driven by the number of pages we track within the reaction time frame. The reaction time not only provides an upper bound on the critical window, but also serves as indicator for the expected number of tracked pages.

In our evaluation, we investigate the number of pages that have to be extracted, and the duration of the observation and search phase. We call the combined duration of both phases the *attack time*.

5.2 Prototype and Test Setup

We implemented our prototype including the functionality required for SEVered based on Kernel-based Virtual Machine (KVM). To start and stop page tracking and change mappings, we extended the KVM API with additional calls, in particular, with *KVM system ioctls* [19]. This allows us to launch the attack from user space by communicating with the KVM kernel module. For page access tracking in KVM we used the technique from [9, 16]. While tracking is active, we record all tracked pages in a list in kernel memory.

Upon the call to stop tracking, KVM returns the list of tracked pages to user space.

We ran KVM on Debian with a page size of 4 KB using an SEV-enabled Linux kernel in version `4.18.13` and QEMU `3.0.50`. We used an AMD EPYC 7251 processor with full support for SEV. We created a victim VM with 2 GB of memory and one of the four available CPU cores. We deployed Apache `2.4.25-3` and nginx `1.10.3-1` for the TLS key scenarios, and OpenSSH `7.4` for the SSH scenario in the VM. The FDE scenario is independent of a service, because the FDE key is a kernel resource exclusively used by the OS. We deployed eleven different web resources on each web server. We used 4096-bit private keys for TLS and SSH and a 256-bit symmetric FDE key for storage encryption with AES-XTS. As target for memory extraction with SEVered, we used a page-sized web resource served by nginx.

To capture the handshake messages for the TLS and SSH key scenarios, we used `tcpdump` with `libpcap`, a library for network packet capturing. For TLS, we captured the *change cipher spec* packet the services send to conclude a TLS handshake (filter `tcp[37] == 0x04`). For SSH, we captured the *new keys* message, which concludes the SSH handshake (filter `tcp[37] == 0x15`). We patched `libpcap` to execute a system call for $Tracking_{Start}$ the moment packet capturing begins, and a call for $Tracking_{End}$ the moment the filtered packet is captured. This tight interconnection minimizes the reaction time. To monitor disk I/O events of the VM in the FDE key scenario, we used the tool `inotifywait` to observe *inotify* events. In particular, the *notify* option allows to detect disk writes on the VM's disk image file. We modified `inotifywait` to issue the calls for $Tracking_{Start}$ right before starting to watch events and for $Tracking_{End}$ as soon as an inotify event is identified.

In real-world scenarios, a tenant's VM can show higher or less activity depending on the load caused by its clients. To simulate this behavior, we executed all our tests based on a *load model* with various *load levels*, representing low to high load. In our model, a load level of nine, for instance, refers to nine requests per second to the VM. We randomly alternate between the services for each request. With a probability of $\frac{300}{301}$, we make a request to one of the resources offered by one of the two web servers with equal probability. With a probability of $\frac{1}{301}$, we initiate an SSH login with a user remaining logged in for two minutes. Compared to the number of web server requests, we execute only few SSH logins, as these usually happen less frequently than requests to a web server. The average duration of the observation phase thus lies in the range of a few seconds to a few hundreds of milliseconds for the web servers and in the range of a few minutes to tens of seconds for SSH. Note that sshd forks a new process for each new SSH connection. When the session terminates, the process exits and purges its SSH key. This means that the search time must be less than two minutes to extract the SSH key before the forked process exits.

We conducted 2,000 independent iterations of our attack for each of the four scenarios on four load levels: level *1*, *9*, *17*, and *25*. We started our attacks at random points in time while the VM processed requests according to the specific load level of our model. As an initial preprocessing step before the search phase, we filtered all execute-accesses. In our scenarios, all secrets are data structures located on non-executable memory pages.
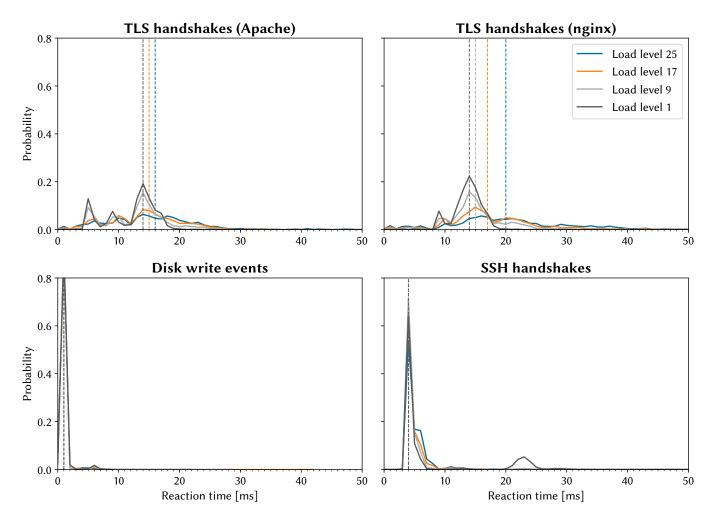
Figure 2: Distribution of the reaction times for all scenarios and load levels. The X-axes show discretized time steps of one millisecond and the Y-axes are normalized to one.

5.3 Success Probability and Reaction Time

In this part, we investigate the success probability and reaction times. The four diagrams in Figure 2 illustrate the distribution of measured reaction times for each scenario. The four graphs in each diagram represent the four load levels. The X-axes are discretized in steps of one millisecond, and the Y-axes are normalized to one. The vertical dashed lines show the median reaction times over all repetitions for each level, providing an upper bound on the median critical window size.

The results for Apache and nginx TLS handshakes are depicted in the top row of Figure 2. Both diagrams show a clear peak for the two lower load levels, indicating a reliable reaction time when the VM is not under high load. For the lowest load level, we can even observe that the reaction time never exceeded 21 milliseconds with Apache and 22 milliseconds with nginx. For higher load levels, more concurrent activities are executed by the VM, and the measurements are more dispersed over time. Consequently, it becomes likely that more pages have to be extracted in the search phase until the secret is found. This led to a maximum reaction time of around

50 milliseconds for both nginx and Apache in rare cases. However, the median reaction time increased to about only 20 milliseconds for nginx, and to about 16 milliseconds for Apache. As the reaction time is an upper bound for the critical window, the latter is smaller than tens of milliseconds for both Apache and nginx. We achieved a very high success rate of around 99.99% for both web servers on all load levels, meaning that we started tracking inside the critical window only in a few cases. The high success rate indicates that the upper bound we measured is a very conservative estimate. This comes from the fact that our prototype requires some time to actually stop tracking and (especially for the TLS scenarios) to recognize $Activity_{End}$. Note that if $Tracking_{Start}$ occurs inside the critical window of a TLS handshake, we still have the chance to observe Use_i of other handshakes being concurrently processed on higher load levels where lots of handshakes are made each second. The critical window can thus be even smaller on higher load levels.

The bottom left diagram in Figure 2 for disk write events shows that our implementation achieved an extremely fast reaction time of one millisecond in the median for each load level. Only in a few

Table 1: Statistics for the median length of the observation and search phases, and for the median number of extracted pages with the median absolute deviation for the different scenarios and load levels.

Use Case	Load Level	Median Page No	MAD Page No	Median Time Observ.	Median Time Search
TLS (nginx)	1	102	5	1.46s	17.72s
	9	116	19	0.37s	15.48s
	17	165	69	0.32s	18.61s
	25	301	160	0.31s	32.71s
TLS (Apache)	1	128	21	1.42s	21.90s
	9	137	40	0.37s	17.95s
	17	154	80	0.33s	17.44s
	25	171	109	0.32s	18.65s
FDE	1	70	8	2.43s	12.24s
	9	71	9	2.15s	9.34s
	17	70	8	2.08s	7.84s
	25	69	9	2.04s	7.37s
SSH	1	7	1	193.36s	1.33s
	9	7	1	27.23s	0.97s
	17	7	1	16.19s	0.83s
	25	7	1	14.41s	0.80s

cases, we encountered a slightly higher reaction time. In contrast to the TLS key scenarios, the behavior was generally independent of the load level. In the TLS key scenarios, the network packets must first be sent by the VM to the network interface, on which the HV executes more time-consuming network packet capturing. The interception of disk write events is less complex and introduces less delay. The success rate for FDE key extraction was about 99.99%, indicating a very small critical window, as confirmed by the upper bound in the graph.

The bottom right diagram in Figure 2 shows that the reaction time for SSH handshakes was four milliseconds in the median and mostly independent of the load level. We encountered only a few samples going up to about 30 milliseconds. As for the TLS scenario, this indicates a small upper bound on the critical window and a possibly quick extraction. Accordingly, our attack had a success rate of 99.98%.

5.4 Attack Time

This part investigates the attack times for each scenario. Table 1 summarizes the relevant statistics for the median number of pages to be extracted and the median attack time for every scenario and load level. For both Apache and nginx, the median number of pages to be extracted until finding the TLS key increased between low and high load levels. We measured an increase of the median from 102 to 301 pages (i.e., 408 to 1,204 KB of memory) for nginx, and from 128 to 171 pages (i.e., 512 to 684 KB of memory) for Apache. Additionally, the Median Absolute Deviation (MAD) increased from 5 to 160 from low to high load for nginx, respectively from 21 to 109

for Apache. The median number of extracted pages was particularly small compared to the median number of total tracked pages, which was for both cases between 1,691 and 2,085 (not listed in Table 1). The median duration of the search phase was between about 15.5 and 32.7 seconds for nginx, and between about 17.5 and 22 seconds for Apache. We measured an average extraction time of around 123 milliseconds for a single page with our SEVered implementation and setup. We measured this time to fluctuate quite frequently in the scale of a few tens of milliseconds. This is why a higher median number of extracted pages did not affect the duration of the search phase in a clearly linear way. The on the fly analysis for a single memory page took about 50 milliseconds. This means that the page extraction performance is the limiting factor of our attack. The higher the load, the less time we required for the observation phase, which ranged from 1.46 to 0.31 seconds in case of nginx, for instance. This is because the probability of quickly observing $Activity_{End}$ increases with a high frequency of requests. To summarize, we measured an attack time between about 16 and 33 seconds in the median for the web services.

For the FDE key scenario, the amount of pages that had to be extracted was very small and mostly independent of the load level. Accordingly, the median number of extracted pages was between 69 and 71 for the different load levels (i.e., 276 to 284 KB of memory) As in the TLS key scenario, this number is small compared to the median number of total tracked pages, which was between 2,526 and 3,433. The overall duration of the search phase was between about 7.4 and 12.3 seconds. The on the fly analysis for a single memory page took only about 2 milliseconds on average. We mostly identified the key as part of the AES key schedule, and only occasionally by the kernel data structure, see Section 4. We measured a slightly decreasing observation time from 2.43 to 2.04 seconds. This indicates that the VM's OS is regularly writing pages to disk, in our case mostly regardless of the load level. In sum, the attack time was between less than 9.5 and 14.7 seconds in the median.

In the SSH scenario, we merely had to extract seven pages in the median with a MAD of one. This is a particularly small number, especially compared to the median number of 10,102 to 11,094 total tracked pages, omitted from Table 1. We measured a median duration of the search phase of about 0.80 to 1.33 seconds. This means that the attack works reliably assuming that the SSH connection lasts at least 1.33 seconds. Similar to the TLS key scenario, the on the fly analysis of a memory page took about 50 milliseconds. With our load model, the observation time of about 14 to 194 seconds was comparably high for the SSH scenario. This is another reason why the number of extracted pages was especially low for the SSH case. In long observation phases, we already tracked a high number of pages before Use_n, making it very unlikely that many pages are tracked within in the reaction time frame at the end of the activity.

5.5 Optimization with Preprocessing

As discussed in Section 3, preprocessing with prioritization and filtering is an optional optimization before the search phase. Preprocessing usually requires a priori knowledge about the use case and behavior of the VM, which may not always be available. This behavior may also vary between different hard- and software configurations. For our evaluation, we already used the knowledge that

all secrets are data structures located on non-executable memory pages. This allowed us to filter execute-accesses from the list of tracked pages. The amount of pages to be extracted was thereby reduced by about 22% on average over all samples.

Use_n was a read-access in 96% of our attacks for the TLS handshakes and in 93% a write-access for the SSH handshakes. For disk write events, Use_n was always a read-access. Whether the page containing the secret is tracked as read- or write-access depends on the other data located within the page. The type of access thus cannot be predetermined with certainty. Filtering of write-accesses could significantly reduce the attack time, but could also reduce the success probability. Also, prioritizing the extraction of read-accesses over write-accesses in the list would boost the attack in most cases, but could also introduce costly outliers.

Another possibility for prioritization is knowledge about the reaction time, as shown in Figure 2. The graphs for the two web server scenarios show that the reaction time was rarely less than eight milliseconds before $Tracking_{End}$. Re-arranging these early accesses further back in the list of tracked pages can thus reduce the amount of pages to be extracted until the secret is found. The same observation can be made for the SSH scenario, where Use_n never happened less than three milliseconds before $Tracking_{End}$. However, in this case the number of pages to be extracted is already so small that further optimization may not be required.

The reaction times in Figure 2 can also help to determine a good criterion for restarting the attack when a secret has not been found after a certain number of extracted pages. For instance, page accesses tracked later than 30 milliseconds before $Tracking_{End}$ are likely exceeding the reaction time frame and thus unlikely to be a candidate for Use_n. This can be used as a criterion to detect that $Tracking_{Start}$ was inside the critical window and to restart the attack from the observation phase.

5.6 Summary

For all evaluated scenarios, both performance indicators are very promising. We found that the critical window was very small in all cases. $Tracking_{Start}$ was thus inside the critical window only in a few cases, resulting in a very high success probability throughout all scenarios and load levels. The most important factor for the attack time, the duration of the search phase, was also very small. Extracting all memory from our VM with 2 GB of main memory would take more than 7 hours with SEVered [16]. Assuming that a key can reside not only on one but on several pages, the naive extraction would require several hours on average to find the key. Our approach can extract secrets faster by several orders of magnitude.

In cases when $Tracking_{Start}$ is inside the critical window, the attack fails and we extract all tracked pages without finding the secret. In such cases, we have to repeat both the observation and search phase. To avoid a lengthy extraction of all tracked pages in unsuccessful attempts, the search can be canceled early when the likelihood of finding the secret drops, according to our evaluated distribution. For the following search phase, all pages extracted in the previous attempts can then be excluded from extraction given that the secret does not change its location. In sum, our results have shown that our prototype is able to quickly and reliably extract different sensitive secrets and performs well even under high load.

6 DISCUSSION

In the following, we discuss further important aspects of our attack and possible countermeasures:

Overhead. The overhead caused by the tracking itself is limited, because each accessed page only triggers a SLAT page fault once. We neither detected perceivable effects like delays in response times in the HV nor inside the VM. The host system and VM remained stable even on the highest load level. We measured only a small additional delay of web and SSH server responses when tracking was active.

Low Memory. When the VM is low on memory, its kernel might try to free memory by swapping out pages, by unmapping file-backed pages, or by killing processes. A page containing the secret might then be re-used by another process or by the kernel during our attack. In such a case, we are still able to extract the memory contents of the page, but its contents might have already been overwritten. We did not encounter such cases in our tests.

Triggering Activities. In our concept, we start tracking at an undefined point in time and do not actively trigger activities to interfere as little as possible with the VM's normal operation. Our concept worked well in our evaluated scenarios, because we extracted frequently used secrets. In the SSH scenario, however, the key may be used rather infrequently. This is, for instance, the case when an administrator logs in to a web server for maintenance only from time to time. When a secret is rarely used but the attacker requires the observation phase to be as short as possible, the attacker can consider the active triggering of an activity. In the SSH scenario, an attacker can actively start a login procedure without a user account. SSH servers use their key for server authentication and wait for the user to authenticate with a default timeout of two minutes. An attacker can thus initiate a login and extract the SSH key before the session timeout without waiting for a legitimate user to login. Note that active triggering might increase the probability of the attack being detected and might not always be possible.

Portability. We expect that our approach can be transferred to other scenarios and configurations than evaluated. Our approach does not depend on specific service or library versions. Furthermore, our approach is not tied to a specific SEV processor and mostly independent of the VM's performance and OS. Our approach can also be leveraged to extract other types of memory, such as confidential code, documents or images. The performance of our approach can differ on systems with other hard- and software configurations. However, we expect the performance to vary only slightly assuming that $Tracking_{End}$ can be observed quickly. We ran several tests in which we assigned our VM more memory, multiple cores and in which we configured the web servers to utilize a high number of worker processes. The performance indicators remained coherent with our evaluation results in all runs.

Countermeasures. A countermeasure against our attack is to prevent the SEVered attack [16], which we rely on for memory extraction. Further, our attack relies on targeted secrets to remain at their memory location during our attack. Purging secrets in memory after use would cause the search phase to fail. We found TLS and FDE keys to always remain at their memory location in

our tests. However, in case of SSH keys, the processes forked by the SSH daemon for initiating new SSH connections purge their private key when a session terminates and then exit. This means that an SSH session must remain open until the secret is extracted. This, for instance, requires a user to remain logged in or a login attempt to remain pending over the time of the search phase, which is less than 1.5 seconds in our case.

Systematically purging all sorts of secrets from main memory after use would require adapting existing software. For some secrets, purging might not be feasible. An example is the FDE key, which is constantly required for disk I/O. A more promising solution is to relocate the most valuable secrets from main memory to dedicated hardware. Since SEVered can only extract contents from main memory, storing secrets in hardware would prevent them from being extracted by a malicious HV. This can, for example, be realized using Hardware Security Modules (HSMs). Additionally, hardware-based disk encryption can be used to protect the FDE key.

7 RELATED WORK

While there are established Virtual Machine Introspection (VMI) frameworks [14, 18, 20] for data analysis and extraction on unencrypted VMs, the systematic extraction of memory from encrypted VMs has not been subject to extensive study. On AMD SEV platforms, the SP protects page encryption and the corresponding keys from the HV. This makes it infeasible to directly read memory contents from SEV-enabled VMs as long as the SP cannot be compromised [13]. Payer [17] early discussed the missing integrity protection on AMD SEV platforms. By remapping memory in the HV, this can be used to extract memory without compromising the SP, as done by the SEVered attack [16]. While SEVered allows the extraction of data, it does not provide concepts for quickly extracting specific secrets.

Buhren et al. presented an attack [11] to gain remote code execution with user privileges on an SEV-enabled VM. Their approach exploits memory remapping to modify the control flow of an SSH service. The first step is an off-line tracing of the system call sequences performed during an SSH login on a comparable, unencrypted VM. The goal of this analysis is to determine the behavior of a VM accessing the login information of the SSH session, the *credentials data structure*. The next step is to wait for a victim user to login to the SSH service. With the information gained in the offline analysis, they identify the memory page containing the user's login information. They then try to illicitly login by remapping the valid user's credential data structure to the one the SSH service creates during the illicit login attempt. This allows the attacker to re-use the victim user's login information. In their evaluation, they achieved a success rate of around 23%. The low rate was primarily caused by the fact that the SSH service may store the *credentials data structure* at different offsets within the page. As a condition for a successful attack, the SSH service must have stored both the victim user's and attacker's credentials data structures at the same page offset. Besides being quite invasive, this approach requires access to a comparable VM, detailed analysis of the SSH service, user interaction, and data being incidentally stored at specific offsets.

The attack described in [8] follows the same goal of gaining remote code execution on an SEV-encrypted VM, but does not exploit remapping. Instead, the authors describe a *ciphertext block move* attack, which also exploits the missing integrity protection. The authors argue that it is possible to move memory contents in physical memory. This is because the HPA is not part of the AES-based encryption scheme itself but is incorporated into the encryption result in a later step with a reversible physical address-based tweak algorithm that uses static parameters. After reversing the tweak, ciphertext can be moved and the tweak re-applied with the target HPA. The authors describe a method that moves the pages to exploit an SSH process. Both the approaches in [11] and [8] were, to the best of our knowledge, not confirmed on real SEV hardware. The *ciphertext block move* attack could possibly be leveraged for the memory extraction as an alternative to the remapping in SEVered.

On the side of defenses, Fidelius [22] is a software-based extension to SEV. This extension is a privileged module separate from the HV that restricts the HV from accessing specific critical resources with *non-bypassable memory isolation*, for instance, to prevent replay attacks. The authors provide a VM lifecycle concept that describes how to start Fidelius and provide tenants proof that the system runs Fidelius in addition to SEV. This requires trusting the Fidelius module instead of the operating HV.

Intel announced the implementation of its own hardware-based memory encryption approach called Multi-Key Total Memory Encryption (MKTME) [12]. According to our understanding of the specification, MKTME does not protect from a malicious or compromised HV, but only from memory attacks from outside. The HV remains, for instance, capable of enabling or disabling the encryption, or to handle the sharing of memory with other VMs.

8 CONCLUSION

We presented an approach for the efficient extraction of secrets from SEV-encrypted VMs. Compared to time-consuming, naive memory extraction, our two-phased approach exfiltrates secrets unobtrusively and quickly with a high success probability. In the first phase, we track the page accesses of an encrypted VM until detecting an event indicating that the VM recently accessed the secret. In the second phase, we leverage an existing attack for memory extraction to systematically retrieve the tracked pages and simultaneously analyze their contents to quickly identify the secret. We presented various use cases for highly sensitive secrets usually found in VMs in cloud scenarios. We performed an evaluation for these cases on a fully SEV-enabled EPYC processor with varying levels of load, usually caused by independent clients not involved in the attack. Our results show that we are able to extract TLS keys after a handshake in less than 15.5 seconds in the median on lower load levels and in no more than about 32.7 seconds in the median on our highest evaluated load level. The extraction of the FDE key after a disk write event took between less than 7.4 seconds and 12.3 seconds in the median. The extraction phase for SSH keys after an SSH handshake took about 0.8 to 1.35 seconds in the median. We expect that our approach can be used for the extraction of further types of secrets, which we are going to investigate in future work.

ACKNOWLEDGMENTS

REFERENCES

[1] Advanced Micro Devices. 2008. Nested Paging. http://developer.amd.com/wordpress/media/2012/10/NPT-WP-1%201-final-TM.pdf.

[2] Advanced Micro Devices. 2018. Secure Encrypted Virtualization API Version 0.16. http://support.amd.com/TechDocs/55766_SEV-KM%20API_Specification.pdf.

[3] Michael Becher, Maximillian Dornseif, and Christian N Klein. 2005. FireWire: All Your Memory Are Belong To Us. *Proceedings of CanSecWest*.

[4] Adam Boileau. 2006. Hit by a bus: Physical access attacks with Firewire. *Presentation, Ruxcon*.

[5] Center for Information Technology Policy at Princeton University. 2008. Memory Research Project Source Code. https://citp.princeton.edu/research/memory/code/.

[6] David Kaplan. 2017. Protecting VM Register State with SEV-ES. White Paper.

[7] Christophe Devine and Guillaume Vissian. 2009. Compromission physique par le bus PCI. *Proceedings of SSTIC*.

[8] Zhao-Hui Du, Zhiwei Ying, Zhenke Ma, Yufei Mai, Phoebe Wang, Jesse Liu, and Jesse Fang. 2017. Secure Encrypted Virtualization is Unsecure. arXiv:cs.CR/1712.05090 https://arxiv.org/abs/1712.05090

[9] Xiao Guangrong. 2016. [PATCH v3 00/11] KVM: x86: Track Guest Page Access. http://www.mail-archive.com/linux-kernel@vger.kernel.org/msg1076006.html.

[10] J. Alex Halderman, Seth D. Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A. Calandrino, Ariel J. Feldman, Jacob Appelbaum, and Edward W. Felten. 2009. Lest We Remember: Cold-boot Attacks on Encryption Keys. *Commun. ACM* 52, 5 (May 2009), 91–98. https://doi.org/10.1145/1506409.1506429

[11] Felicitas Hetzelt and Robert Buhren. 2017. Security Analysis of Encrypted Virtual Machines. In *Proceedings of the 13th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments (VEE '17)*. ACM, New York, NY, USA, 129–142. https://doi.org/10.1145/3050748.3050763

[12] Intel. 2017. Intel Architecture Memory Encryption Technologies Specification. https://software.intel.com/sites/default/files/managed/a5/16/Multi-Key-Total-Memory-Encryption-Spec.pdf.

[13] CTS Labs. 2018. *Severe Security Advisory on AMD Processors*. Technical Report.

[14] LibVMI Project. 2015. LibVMI Virtual Machine Introspection. http://libvmi.com/.

[15] Microsoft. 2017. Microsoft Security Bulletin MS17-008 - Critical. https://technet.microsoft.com/en-us/library/security/ms17-008.aspx.

[16] Mathias Morbitzer, Manuel Huber, Julian Horsch, and Sascha Wessel. 2018. SEVered: Subverting AMD's Virtual Machine Encryption. In *Proceedings of the 11th European Workshop on Systems Security (EuroSec'18)*. ACM, New York, NY, USA, Article 1, 6 pages. https://doi.org/10.1145/3193111.3193112

[17] Mathias Payer. 2016. AMD SEV Attack Surface: a Tale of too Much Trust. https://nebelwelt.net/blog/20160922-AMD-SEV-attack-surface.html.

[18] Rekall Forensics. 2018. Rekall. http://www.rekall-forensic.com/.

[19] The Linux Kernel Organization. 2018. The Definitive KVM (Kernel-based Virtual Machine) API Documentation. https://www.kernel.org/doc/Documentation/virtual/kvm/api.txt.

[20] The Volatility Foundation. 2018. Open Source Memory Forensics. https://www.volatilityfoundation.org/.

[21] VMware. 2017. VMSA-2017-0006: VMware ESXi, Workstation and Fusion Updates Address Critical and Moderate Security Issues. https://www.vmware.com/security/advisories/VMSA-2017-0006.html.

[22] Yuming Wu, Yutao Liu, Ruifeng Liu, Haibo Chen, Binyu Zang, and Haibing Guan. 2018. Comprehensive VM Protection Against Untrusted Hypervisor Through Retrofitted AMD Memory Encryption. In *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. 441–453. https://doi.org/10.1109/HPCA.2018.00045

[23] Xenproject.org Security Team. 2017. x86: Broken Check in memory_exchange() Permits PV Guest Breakout. https://xenbits.xen.org/xsa/advisory-212.html.