

TME-Box: Scalable In-Process Isolation through Intel TME-MK Memory Encryption

Martin Unterguggenberger*, Lukas Lamster*, David Schrammel*, Martin Schwarzl†, Stefan Mangard*

*Graz University of Technology: {firstname.lastname}@iaik.tugraz.at

†Cloudflare, Inc.: mschwarzl@cloudflare.com

Abstract—Efficient cloud computing relies on in-process isolation to optimize performance by running workloads within a single process. Without heavy-weight process isolation, memory safety errors pose a significant security threat by allowing an adversary to extract or corrupt the private data of other co-located tenants. Existing in-process isolation mechanisms are not suitable for modern cloud requirements, e.g., MPK’s 16 protection domains are insufficient to isolate thousands of cloud workers per process. Consequently, cloud service providers have a strong need for lightweight in-process isolation on commodity x86 machines.

This paper presents TME-Box, a novel isolation technique that enables fine-grained and scalable sandboxing on commodity x86 CPUs. By repurposing Intel TME-MK, which is intended for the encryption of virtual machines, TME-Box offers lightweight and efficient in-process isolation. TME-Box enforces that sandboxes use their designated encryption keys for memory interactions through compiler instrumentation. This cryptographic isolation enables fine-grained access control, from single cache lines to full pages, and supports flexible data relocation. In addition, the design of TME-Box allows the efficient isolation of up to 32K concurrent sandboxes. We present a performance-optimized TME-Box prototype, utilizing x86 segment-based addressing, that showcases geometric performance overheads of 5.2 % for data isolation and 9.7 % for code and data isolation, evaluated with the SPEC CPU2017 benchmark suite.

I. INTRODUCTION

Cloud computing must be highly optimized for performance and efficiency, processing requests with the lowest possible latencies. Various techniques and optimizations are implemented at the software and infrastructural level of the cloud architecture. Common optimizations are the exclusion of protection mechanisms, like process isolation, in favor of low start-up times and fast execution for serverless applications in a multi-tenant environment, e.g., function-as-a-service (FaaS) applications [16], [20], [22]. However, omitting heavy-weight protection mechanisms such as process isolation introduces substantial security risks. A single memory safety vulnerability [54], [82] can result in a full compromise of the entire cloud computing system. Consequently, attackers can compromise other tenant’s highly sensitive data.

Security vulnerabilities, such as Heartbleed [21] and Cloudbleed [28], [63], exemplify the severity of this attack surface. In particular, Heartbleed allowed a remote adversary to leak private data (e.g., confidential key material) by exploiting a buffer-overread error due to improper input sanitization triggered via maliciously crafted network packets. Without the strong isolation of memory resources, exploitable memory safety errors enable the leakage or corruption of sensitive process memory, such as private keys or authentication tokens. To address this security threat while preserving the level of optimization, cloud service providers are replacing process isolation with in-process sandboxes [17], e.g., by using sandboxed languages like JavaScript [29], to efficiently support a high number of tenants per machine. It is evident that cloud service providers have a strong need for scalable and efficient in-process isolation mechanisms that segregate memory resources of cloud-co-located tenants.

Common isolation techniques harden software systems by applying code instrumentation for logical protection with dynamic runtime checks. Software-based fault isolation (SFI) [81], [85] is an effective technique that instruments memory and control-flow operations to keep them within predefined regions. Language-level isolation relies on partitioning the virtual address space to separate process memory [27], [29]. This procedure enables software sandboxing, effectively restricting access for the sandbox to its assigned memory region. Due to its efficiency, SFI has also been used in practice, as seen in the Google Native Client (NaCl) [73], [89] sandbox. However, SFI can only offer coarse-grained sandboxing by isolating predetermined and contiguous memory regions without interleaving. This means that SFI lacks fine-grained, object-level access control and the ability to dynamically manage resources in memory.

The efficiency of isolation can be significantly enhanced with hardware support. Hardware-based isolation mechanisms, such as protection keys for user space (PKU) [61], enable the restriction of access to memory at a page granularity. Specifically, memory protection keys (MPK) use a 4-bit protection key located in the page table entries (PTE) to assign and enforce access policies for the respective page. These policies are software-controlled via a processor register, associating two bits for every protection key to dynamically control read and write permissions. Thereby, MPK allows the partitioning of the process memory into 16 distinct regions, often called domains, and selectively controls their access.

Various PKU-based protection schemes [8], [31], [71], [72], [84] are proposed to establish lightweight and transparent in-process isolation that is efficiently enforced in hardware. However, MPK is limited by its 4-bit key size, only offering 16 distinct protection domains. This is restrictive for software systems that may require thousands of concurrent domains, e.g., Cloudflare Workers [16], [17], within a single process. Also, MPK is bound to the page granularity and cannot provide sub-page granular isolation.

In-process isolation is a highly demanded security feature for modern computing systems, such as cloud computing, imposing hard requirements on sandboxing techniques. The sandbox must provide scalable isolation with fine-grained access control applicable to individual objects and memory pages. Additionally, the technique must support thousands of concurrent sandboxes, e.g., to isolate cloud workers, and must be available on commodity x86 server-class processors.

Contributions. In this paper, we present TME-Box, a novel sandboxing technique for scalable in-process isolation that repurposes Intel’s newly introduced total memory encryption multi-key (TME-MK) [32] technology. TME-MK’s runtime memory encryption is essentially used by the Intel trusted domain extensions (TDX) [15], [33], their confidential computing architecture, to ensure the confidentiality and integrity of memory resources. This encryption and other architectural elements (e.g., secure extended page tables (EPT) [36]) allow TDX to provide strong isolation of virtual machines (VMs), which run isolated workloads on the same physical machine.

Intel TME-MK provides page-granular encryption of the computer’s physical memory, which is primarily designed for the cryptographic isolation of VMs. This enables hypervisors to encrypt VMs and containers [34], as well as to protect against physical attacks on the memory subsystem [30], [33]. However, the heavy-weight VM isolation of TDX incurs significant performance overhead, and TME-MK’s encryption has not yet been applied for generic and efficient in-process memory isolation comparable to MPK.

TME-Box extends the application of the TME-MK memory encryption beyond hardware-isolated VMs by providing fine-grained and scalable in-process isolation on commodity x86 CPUs. TME-Box deploys compiler instrumentation to isolate the code and data of mutually untrusted sandboxes. Specifically, we enforce the usage of sandbox-specific TME-MK encryption keys by controlling the base address and index of memory operations. Thereby, memory resources of the in-process sandbox are cryptographically isolated, leading to the detection of unauthorized memory access through hardware-backed integrity protection. Repurposing the TME-MK memory encryption for software sandboxing offers several advantages over established in-process isolation mechanisms and addresses the hard requirements of modern computing systems, such as cloud computing.

First, TME-Box enforces fine-grained access control of memory resources. Through the use of page aliasing, TME-Box achieves sub-page granular encryption and, thus, fine-grained memory isolation. Particularly, our design allows for

scalable isolation granularities ranging from single cache lines to full pages, especially relevant in modern cloud settings. In contrast, SFI-based mechanisms can only isolate coarse-grained, predetermined, and contiguous memory regions, restricting access by partitioning the virtual address space.

Second, TME-Box supports the flexible relocation of data in memory, which is particularly important for data-centric computation. Data relocation is essential in the cloud, e.g., when allocator caches with smaller heap slots are used. This enables flexible memory management for the allocator, providing continuous memory utilization across different sandboxes and efficient memory migration.

Third, our TME-Box design leverages a larger amount of key identifiers (keyIDs), enabling the support of thousands of concurrent sandboxes. Compared to MPK’s 4-bit protection keys (enabling 16 distinct isolation domains), our design leverages TME-MK that supports up to 15-bit keyIDs, addressing up to 32K encryption keys [32] that we repurpose to cryptographically isolate sandboxes. This is particularly relevant for software systems that require a large number of sandboxes, e.g., cloud workers [16], [17], within a single process. Additionally, our design supports frequent policy changes from user space without kernel interaction.

Complementary to our design, we detail our prototype implementation of the TME-Box framework, consisting of an LLVM compiler toolchain and a security-hardened allocator. Furthermore, we outline architecture-specific optimizations, such as x86 segment-based addressing [39], to achieve practical performance results for SPEC CPU2017 [10] and NGINX. We showcase a geomean overhead of 5.2% for data isolation and 9.7% for code and data isolation using SPEC CPU2017.

In summary, the main contributions of this work are:

- **TME-Box.** We are the first to present a novel in-process isolation technique by repurposing the Intel TME-MK memory encryption on commodity x86 CPUs.
- **New Insights on Sandboxing.** We provide new insights on hardware-assisted sandboxing through memory encryption for modern cloud settings, enabling fine-grained and scalable isolation from single cache lines to full pages while supporting up to 32K sandboxes.
- **Prototype and Evaluation.** We implement an optimized prototype of TME-Box, utilizing x86 segment-based addressing, that showcases competitive performance results for SPEC CPU2017 and NGINX.
- **Security Analysis.** We systematically analyze the security threats and outline the derived security properties of our sandbox design.

Outline. The paper is structured as follows. Section II provides the background of this work. Section III defines our threat model. Section IV and Section V describe TME-Box’s design and implementation. Section VI and Section VII provide the security analysis and evaluation. Section VIII discusses related work, and Section IX concludes this paper.

II. BACKGROUND

This section provides the background on address translation with virtual memory, software-based fault isolation (SFI), and Intel defensive execution technologies.

A. Address Translation

In modern operating systems, virtual memory organizes memory into contiguous blocks known as *memory pages*, which are typically 4kB in size on the x86-64 architecture. Pages are managed in a hierarchical structure called *page tables*, enabling each user space application to have a distinct virtual representation of the computer’s physical memory.

Address translation in paging systems is performed by the memory management unit (MMU), which is responsible for resolving the virtual-to-physical mappings. When accessing memory, the CPU translates a virtual address to a physical address via a page table walk. The MMU uses the corresponding page table entries (PTE) that contain the physical page number (PPN) (and associated access permissions) to resolve the virtual address. Frequently used virtual-to-physical mappings are cached in the translation look-aside buffer (TLB) to enhance system performance.

Paging provides process isolation through virtual memory as an abstraction of the computer’s physical memory. Access protection is enforced in hardware and managed by the operating system, thus preventing processes from illegally accessing other processes’ memory.

Typically, modern processors provide support for 48-bit or 57-bit virtual address spaces used for the address translation. The virtualization of the computer’s physical memory also allows multiple virtual addresses to refer to the same physical memory. This is called *aliasing* and is often used for shared memory (*i.e.*, code and data) across different processes.

B. Software-based Fault Isolation

Memory safety vulnerabilities [80] frequently occur in complex software written in memory-unsafe programming languages (e.g., C, C++) [54], [82]. This is critical for software security, as an adversary can exploit memory safety vulnerabilities to compromise the target system [2], [65]. For instance, out-of-bounds (OOB) errors, such as buffer over-reads and over-writes, enable an adversary to illegally access resources in memory. The exploitation of memory safety errors allows the attacker to leak or corrupt sensitive data.

To mitigate this security threat, software sandboxing is used. Sandboxes aim to reduce the attack surface of memory errors by isolating individual parts of the software. Software-based fault isolation (SFI) [85] is a defense technique that implements in-process isolation (also referred to as intra-process or sub-process isolation), separating memory resources for individual software components within a single process. Each SFI sandbox consists of an isolated data region containing runtime data (e.g., stack and heap memory) and an isolated code region where the code of the sandbox resides [81]. Consequently, the SFI sandbox needs to restrict data access and control-flow transfers. Thus, all memory operations (*i.e.*,

read and write accesses) performed by the sandbox’s code are only allowed to access data within this isolated data region. In addition, control-flow transfers must remain within the sandbox’s code region (or target call sites that correspond to trusted runtime calls) [81].

Data protection is achieved in traditional SFI systems by either using address checking or address masking, which instruments all memory operations (typically performed by the compiler) and enforces memory references to stay inside the sandbox boundaries [51], [85]. Modern SFI approaches implement data protection via the introduction of guard zones [90]. Guard zones refer to unmapped memory regions where memory accesses result in a page fault, thereby detecting accesses outside the sandbox. SFI systems using guard zones need to ensure that memory operations remain within the sandbox’s data and guard region by controlling the base and index register of memory operations. For instance, Google Native Client (NaCl) [73], [89] uses 4 GB virtual memory regions surrounded by 40 GB guard zones to provide software sandboxing.

Control-flow transfers are restricted by SFI, enforcing that the sandboxed code remains in the designated code region [73], [89]. This is achieved by instrumenting indirect control-flow transfers and function returns. Therefore, SFI prevents memory safety errors from corrupting memory outside a vulnerable software component, protecting the remaining system from exploitation and reducing the attack surface in the scenario of memory attacks. The concept of SFI has been revisited for modern CPU architectures, such as x86 and ARM, proposing architecture-specific optimizations [23], [55], [73], [88], [89], [93].

C. Intel Defensive Execution Technologies

The Intel x86 architecture offers platform-specific hardware features that can be used for runtime protection.

Intel Control-Flow Enforcement Technology. Control-flow enforcement technology (CET) [79] is a set of processor extensions designed to implement control-flow integrity (CFI) [1], [12], [13] measures into the CPU hardware. CFI can be loosely categorized into two types: forward-edge and backward-edge CFI, which aim to protect function pointers and return addresses, respectively.

In the context of forward-edge CFI, Intel CET integrates indirect branch tracking (IBT), utilizing so-called landing pads identified by the `endbr64` instructions inserted into function entries by the compiler. Valid indirect jump targets are limited to function entries, thus reducing the attack surface of control-flow hijacking attacks such as jump-oriented programming (JOP) [9]. In terms of backward-edge CFI, Intel CET integrates a hardware shadow stack feature, which applies protection for return addresses. The shadow stack feature achieves this by copying the return address onto the shadow stack, which is inaccessible to an attacker when entering a function. On function exit, the return address restored from the regular stack is compared against the one on the shadow stack. Mismatches are detected in hardware, and the potential

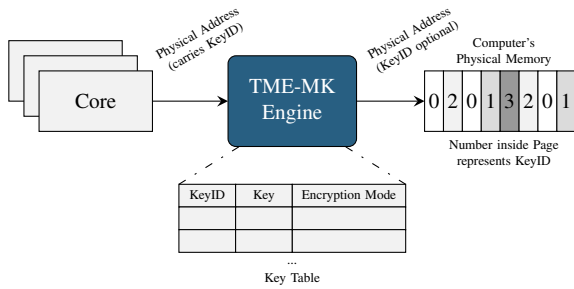


Fig. 1: High-level overview of the TME-MK encryption engine. The encryption engine processes the transferred data depending on the key identifier (keyID) used, which is encoded in the upper bits of the physical address. TME-MK uses a dedicated key table to resolve mappings from keyIDs to encryption keys and encryption modes.

corruption of return addresses is mitigated. Thereby, CET thwarts code-reuse attacks such as return-into-libc [77] and return-oriented programming (ROP) [11].

Intel Memory Protection Keys. Memory protection keys (MPK) [61] are a hardware feature that enforces page-level memory protection and allows changing permissions without requiring page table modifications.

MPK adds a 4-bit protection key encoded into the page table entries (PTEs) and extends the processor architecture with a user space protection key register (PKRU). The PKRU register enables the enforcement of software-controlled read and write permissions of pages by associating two bits for every distinct protection key. MPK performs logical integrity checks during the address translation in the MMU by comparing the protection key with the current access permissions defined by the PKRU register. In this way, MPK allows partitioning process memory into 16 distinct memory regions, dynamically controlling access permissions. The PKRU register is accessible in user mode, which increases the performance of policy changes since they do not require kernel interaction.

While originally introduced to enable execute-only memory and the locking-away of secrets, MPK has also been used to facilitate in-process isolation of untrusted code [31], [71], [84]. However, using MPK for isolation has some drawbacks. For instance, MPK’s protection is limited for certain software architectures due to the relatively small 4-bit protection key. Particularly, this means that MPK can only offer 16 distinct domains, which is restrictive for applications that may require thousands of separate domains (e.g., FaaS applications) to isolate the memory of a single process. Moreover, MPK applies its protection at the page level. Thus, fine-grained (*i.e.*, sub-page granular) isolation cannot be provided.

Intel Total Memory Encryption Multi-Key. Total memory encryption (TME) [32] is Intel’s memory encryption technology that enables transparent encryption of DRAM data with a single encryption key.

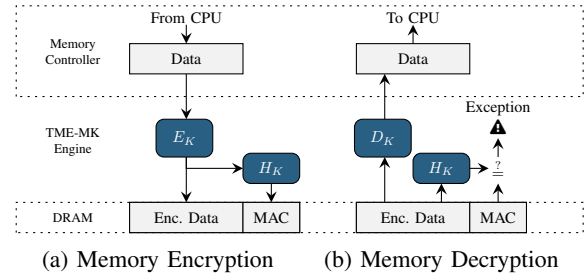


Fig. 2: Overview of TME-MK’s memory encryption with integrity support. Memory operations leverage a specific encryption key that is determined by the keyID to encrypt and decrypt data located in DRAM. Memory accesses using an incorrect keyID are detected by a MAC mismatch and result in a hardware exception.

The total memory encryption multi-key (TME-MK)¹ extension enhances Intel TME with support for multiple encryption keys. In general, memory encryption is a widely used technology that can provide confidentiality and integrity of DRAM memory. The TME-MK feature is currently used for the cryptographic isolation of virtual machines [34] and for protection against physical attacks, such as cold boot attacks [30], [33].

Figure 1 provides a high-level overview of the Intel TME-MK encryption engine. The TME-MK engine is located between the CPU core and the memory controller that interacts with the DRAM memory. Moreover, TME-MK allows the use of up to 2^{15} keys, and it is platform-dependent how many are actually implemented. The encryption engine maintains a key table that manages mappings of different key identifiers (keyIDs) that correspond to encryption keys (and encryption modes). Memory operations have a specific keyID encoded into the upper part of the request’s physical address, allowing the encryption of memory pages with different keys. Note that the physical address and keyID of the memory page are stored in the corresponding page table entry (PTE) managed by the operating system.

Figure 2 illustrates the encryption procedure of TME-MK with integrity support. The TME-MK encryption engine uses AES [18], [19] in XTS mode [50], [67], supporting both 128-bit and 256-bit encryption keys. In addition, with the introduction of the Intel trust domain extensions (TDX) [15], [33], authenticated encryption is integrated. This extends TME-MK with support for data integrity using a message authentication code (MAC). The MAC is computed using the SHA-3 [7] secure hash algorithm. Memory operations are performed using the associated keyID that selects the used encryption key. When writing to memory, TME-MK encrypts the data and computes the MAC to provide confidentiality and integrity of DRAM data (cf. Figure 2a - Memory Encryption).

¹Intel Total Memory Encryption Multi-Key (TME-MK) [32] extension was introduced with the 3rd generation of Intel Xeon Scalable server processors.

When reading encrypted data from DRAM, it decrypts and authenticates the data by recomputing and comparing the associated MAC. An integrity violation is detected within the cryptographic bounds of the MAC (cf. Figure 2b - Memory Decryption). In case of an integrity violation, either a fixed pattern is returned to prevent ciphertext analysis [33], or a hardware exception is raised [70].

TME-MK also allows the configuration of the used encryption mode, *i.e.*, the key size, and to enable integrity support. When using integrity, each MAC covers one cache line. This means that the smallest usable granularity for encryption with integrity mode is also cache line-sized (*i.e.*, 64 B).

III. THREAT MODEL

We consider a strong adversary that intends to exploit one or several memory safety vulnerabilities [80] present in the target unprivileged user space program, thereby gaining unauthorized access to resources in memory. The attacker obtains an arbitrary (read and write) memory access primitive, *e.g.*, by triggering a buffer over-read or over-write error via maliciously crafted user input. They then exploit this vulnerability in an attempt to leak or corrupt sensitive process memory (*e.g.*, confidential data such as private keys or authentication tokens).

Also, we assume that the adversary possesses knowledge of the process’s address space layout (*i.e.*, the attacker can circumvent ASLR [26], [78]). However, we assume that privileged software (*i.e.*, the operating system or hypervisor) is benign and free of exploitable programming errors. Moreover, we presume that Write-XOR-Execute is enabled by default. Thus, the attacker cannot perform code-injection attacks [6] due to the enforcement of the no-execution policy on writable memory. Side channels [41], [48] and fault injection attacks [40], [56] are considered out-of-scope for this work.

IV. TME-BOX SYSTEM DESIGN

This section presents TME-Box, a novel technique for strong and efficient in-process isolation built on top of the hardware-backed TME-MK encryption. TME-Box repurposes Intel TME-MK, intended to encrypt the memory of virtual machines, for fine-grained and scalable in-process isolation on commodity Intel x86 CPUs. By applying compiler instrumentation, TME-Box enforces that sandboxes use a designated encryption key for memory interactions. Thereby, TME-Box cryptographically isolates memory, allowing the detection of unauthorized accesses from mutually untrusted sandboxes.

A. Secure System Architecture

Design Properties. In the following, we outline the main design properties of TME-Box, addressing the hard requirements of modern computing systems, such as cloud computing. We enable scalable isolation and flexible data relocation without introducing new software dependencies, *e.g.*, hard constraints for the memory allocator. We provide cryptographic in-process isolation, thereby supporting a large number of sandboxes (*e.g.*, for isolating cloud workers) and detecting unauthorized

access through integrity exceptions. In this way, TME-Box enables hardware-supported in-process isolation on commodity x86 machines for modern cloud settings.

- **Scalable isolation:** TME-Box enforces fine-grained and scalable access control, supporting isolation granularities ranging from individual cache lines to full pages.
- **Flexible data relocation:** TME-Box allows flexible relocation of data in memory, enabling efficient memory migration and continuous memory operation.
- **Number of sandboxes:** TME-Box supports up to 32K sandboxes, as Intel TME-MK provides up to 2^{15} encryption keys that we repurpose for cryptographic isolation.
- **Integrity enforcement:** The TME-MK encryption engine ensures data integrity for memory operations, allowing us to detect unauthorized sandbox accesses.
- **Commodity hardware:** TME-Box repurposes Intel TME-MK, a platform-specific hardware feature available on commodity x86-64 CPUs.

Overview. At its core, TME-Box repurposes the TME-MK encryption engine for efficient and secure in-process isolation by enforcing the usage of a sandbox-specific alias that represents a TME-MK key identifier (keyID). Each sandbox is assigned its dedicated keyID, mapping to the sandbox’s encryption key. Moreover, compiler-based code instrumentation ensures that all memory interactions of the sandbox use the sandbox-specific keyID, thereby cryptographically isolating the memory of sandboxes from each other. To achieve this, TME-Box stores the sandbox’s base address in a separate CPU register, protecting it from malicious access. This register can either be a general-purpose register or an x86 segment register [39]. TME-Box enforces memory isolation by controlling the base address and index of memory operations.

TME-Box leverages different views on the computer’s physical memory to provide scalable isolation for sandboxes. The base address is used to select a sandbox-specific keyID through the associated memory alias, ensuring that memory accesses within the sandbox always use its respective encryption key. If a sandbox attempts to access data from another sandbox with an incorrect keyID, the underlying encryption triggers an exception, detecting unauthorized access.

Thereby, TME-Box achieves isolation granularities ranging from individual cache lines to full pages. We combine flexible and software-controlled access policies with the cryptographic integrity enforcement of TME-MK. This means that we are not limited to contiguous memory ranges and can quickly re-assign individual small memory granules to different sandboxes. Sandboxes can grow or shrink in size, down to the granularity of the underlying memory encryption, which is 64 B cache line size. TME-Box also allows the flexible relocation of memory depending on runtime constraints. Moreover, we can provide numerous sandboxes, *i.e.*, TME-MK supports up to 2^{15} keyIDs addressing 32K encryption keys that we use for isolation.

TME-Box ensures the isolation of runtime data (*e.g.*, stack and heap memory) and restricts control-flow transfers of its code region. During the startup of each sandbox, it initializes

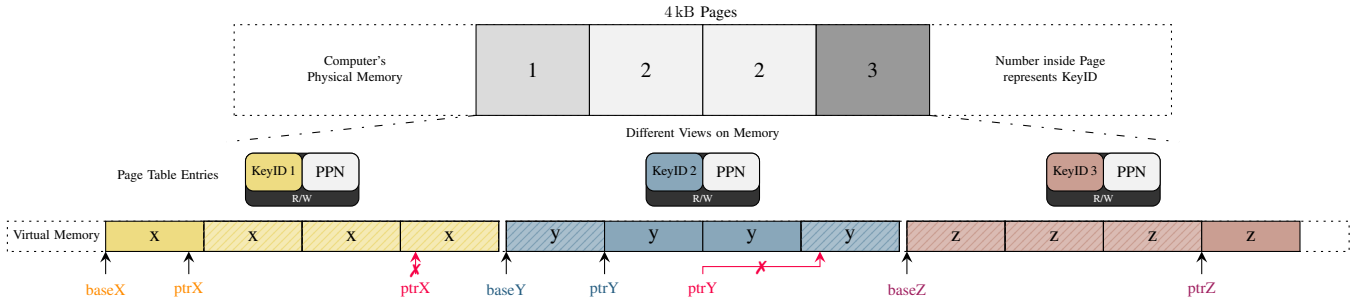


Fig. 3: Overview of TME-Box’s page granular isolation by illustrating the virtual address space of three different sandboxes that are mapped to four physical memory pages. Memory operations of the sandboxes x , y , and z use the keyIDs 1, 2, and 3, respectively. The physical pages are encrypted with the encryption keys of their corresponding sandbox, restricting access solely to their keyIDs. When a sandbox accesses pages that are not assigned to it, *i.e.*, unauthorized access of a non-adjacent or adjacent page, the violation is detected by the integrity checks of the TME-MK encryption engine.

the sandbox’s stack and static memory with the corresponding keyID and sets up the base address register. The virtual address space is mapped so that required memory regions, such as heap memory, are accessible through each sandbox’s designated keyID. For this, we use memory aliasing with different keyIDs so that multiple sandboxes can share the same underlying physical memory. This approach provides scalable isolation granularities for dynamic memory and enables flexible resource management. Additionally, TME-Box instruments control-flow transfers to ensure that they reside within the sandbox’s code region. Our compiler instruments indirect function calls and jumps, *i.e.*, function pointers, by applying address masking. Similarly, return addresses are instrumented to ensure control-flow transfers remain within the sandbox.

The memory allocator can dynamically adjust security policies for heap memory, *i.e.*, enable and revoke memory access. The TME-MK hardware performs integrity checks, providing the ability to easily manage and relocate data in memory. Once the keyIDs are set for the virtual address space, individual cache lines of memory can be assigned to a sandbox with a single memory write without kernel interaction. This enables efficient resource management through the flexible relocation of data. Our approach enhances performance, especially when frequent policy changes occur, applying fast changes to the protected regions of the sandbox’s memory.

B. Scalable Memory Isolation and Flexible Data Relocation

In the following, we discuss TME-Box’s scalable isolation and flexible data relocation in memory. We first detail how entire memory pages (or ranges of pages) are isolated from each other. Next, we elaborate on the fine-grained memory isolation and flexible relocation that TME-Box offers.

Scalable Memory Isolation. TME-Box can provide isolation for full pages assigned to different sandboxes. Our software design maps the virtual address space of the sandboxes to the same physical memory through page aliasing. Furthermore, our design leverages compiler instrumentation to enforce that the memory operations of the sandbox always use its associated keyID located in the PTE. By controlling the base address

of memory operations, we ensure that accesses are performed with the corresponding encryption key of the sandbox and validated by TME-MK.

Moreover, this allows the memory allocator to efficiently manage the memory for different sandboxes by initializing the distinct memory locations (*i.e.*, writing an entire cache line with a keyID to initialize the MAC). TME-Box supports the isolation of a large number of pages, including non-contiguous memory regions. Figure 3 illustrates TME-Box’s page-granular isolation, showing how different sandboxes can have distinct views of the computer’s physical memory.

Specifically, the virtual address space of three distinct sandboxes (x , y , and z) is mapped to four pages of the computer’s physical memory. The sandboxes use the keyIDs 1, 2, and 3, respectively, to cryptographically isolate their memory. Each sandbox is restricted to accessing its own memory since the physical pages are encrypted with different encryption keys of the corresponding sandbox, *i.e.*, memory accesses are only granted by using the correct sandbox-specific encryption key. The example shows unauthorized access to a non-adjacent page by sandbox x and to an adjacent page by sandbox y . Invalid accesses are detected by TME-MK’s cryptographic integrity checks. This allows us to assign and isolate individual pages for different sandboxes.

In addition to page granular isolation, our design also enables more fine-grained access control on the level of individual cache lines within a single page. This sub-page granular isolation is achieved by page aliasing, where two (or more) sandboxes, with different keyIDs each, have memory assigned on the same physical page.

Figure 4 illustrates the sub-page granular isolation of TME-Box, allowing fine-grained resource management for distinct objects associated with a sandbox. In the depicted example, two sandboxes, y and z , use parts of the memory of the same physical page. The sandboxes use different keyIDs, and thus, their stored data is encrypted using different TME-MK encryption keys. This example shows that even when multiple sandboxes store their own respective data co-located on a

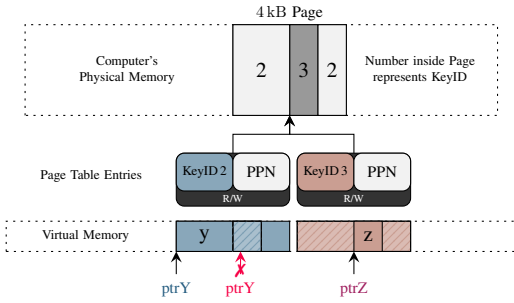


Fig. 4: Overview of TME-Box’s sub-page granular isolation. Aliasing allows the fine-grained encryption of the sandbox’s memory, which is applicable to individual objects located on the same page. Here, two sandboxes share the same physical page to store their private data. Each sandbox can access memory only with its assigned keyID through the correct alias. Hence, parts of the physical page are encrypted differently, with the respective sandbox’s encryption key. This limits access for a sandbox to their respective part of the physical page since accessing non-owned data (with the wrong keyID) is detected by the hardware-based integrity checks.

physical page, the data is still isolated from each other. TME-Box’s fine-grained isolation can be applied to individual cache lines, as this is the smallest granularity of TME-MK with integrity support. This means we can isolate memory without constraining memory allocation of the process, e.g., by relying on a page granularity.

Flexible Relocation of Data. Our scheme enables the flexible relocation of data located in memory, which is important for software systems that focus on data-centric computation. This allows the memory allocator to efficiently migrate memory, thereby reducing the overall memory fragmentation and consumption. Figure 5 provides an overview of TME-Box data relocation. The figure shows the relocation of sparsely allocated cache lines from different sandboxes on memory pages A and B to a single page, page C. This allows the usage of allocator caches and results in a more efficient management of fragmented memory resources.

In contrast, the partitioning of the virtual address space only allows to operate on contiguous memory regions, *i.e.*, on the granularity of a large number of pages. Such systems cannot migrate allocations from different sandboxes to the same page. Here, TME-Box’s approach offers an advantage by allowing fine-grained and flexible relocation of data across different sandboxes.

C. Isolation of In-Process Sandboxes

TME-Box isolates code and data for sandboxing. This includes the fine-grained isolation of data in memory, restricting access for mutually untrusted sandboxes. Additionally, TME-Box ensures the isolated execution of sandboxed code.

Integrity Enforcement. TME-Box repurposes Intel TME-MK to achieve scalable memory isolation. Any violation of TME-Box’s memory access policies, *i.e.*, one sandbox

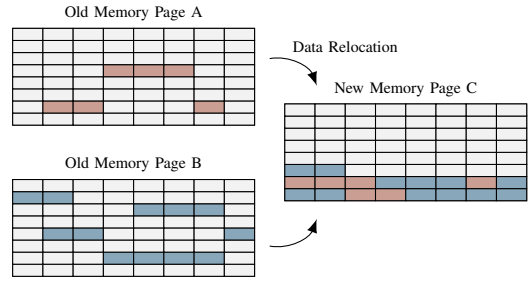


Fig. 5: Overview of TME-Box’s flexible data relocation. This allows for efficient memory migration of sparsely allocated cache lines to different memory locations. When reallocating memory objects, the allocator can move allocations from different sandboxes to the same physical page. E.g., when a page only contains a single object, an allocator can decide to move the object to another already-used page, such that the old page becomes entirely unused. Then, the now-unused old physical page can be freed and reclaimed for other tasks.

trying to access another sandbox’s memory, is detected by the TME-MK encryption engine. TME-Box relies on TME-MK’s integrity checks to detect sandbox access violations. In particular, TME-MK associates a cryptographic MAC that is 28-bit in size with each 64 B cache line [33].

Memory is initialized by writing the entire cache line, which sets up the MAC with the used keyID, thus providing data integrity. For subsequent memory reads, TME-MK authenticates the MAC using the stored DRAM data and keyID of the access. TME-Box leverages this authentication procedure to detect unauthorized memory accesses of sandboxes.

Memory requests with an incorrect keyID result in a MAC mismatch, triggering an exception. Similarly, (partial) writes to a cache line with an incorrect keyID corrupts the MAC associated with the cache line. In this case, the access violation is detected when the sandbox owning the memory location performs a read access. This allows TME-Box to detect access violations within the security bounds provided by the cryptographic MAC.

Isolated Execution. In addition to cryptographically isolating memory resources, software sandboxing requires isolating the executed code, *i.e.*, control-flow transfers remain within a sandbox’s isolated code region. TME-Box achieves this through compiler instrumentation, which modifies forward-edge and backward-edge control-flow transfers to enforce this property.

For forward-edge transfers, TME-Box instruments indirect function calls and jumps (*i.e.*, function pointer dereferences) to stay within the sandbox’s isolated code region. Address masking is applied to these function pointers, enforcing that code execution is confined within the sandbox.

Similarly, backward-edge transfers, such as function returns, are transformed and instrumented to ensure that return addresses target code within the sandbox. Additionally, direct function calls are analyzed during compile-time to ensure that

they target valid call sites, *i.e.*, authorized call targets within the sandbox or trusted runtime calls.

Alternatively, Intel CET [79] can be used to protect control-flow transfers in the TME-Box sandbox. The CET shadow stack feature can replace software-based return address instrumentation for backward-edge CFI, thus enhancing security and performance by protecting return addresses in hardware. In terms of forward-edge CFI, the IBT feature of CET would also provide an additional layer of security. Although indirect function calls would still need to be instrumented to reside within the sandbox code region, the jump targets would be limited to function entries, thus increasing security.

V. PROTOTYPE IMPLEMENTATION

In this section, we detail our prototype implementation. The TME-Box framework comprises an LLVM [45] compiler extension and a security-hardened memory allocator. In addition, we use a Linux kernel patch that enables us to control the TME-MK encryption engine.

A. Compiler Extension

We base our prototype implementation on the LLVM compiler infrastructure [45] (version 14.0.0). Our compiler extension consists of a set of individual compiler passes required for TME-Box’s isolation of code and data. Moreover, the TME-Box framework is fully parameterizable regarding the utilized virtual address space (e.g., 57-bit virtual addressing).

We also detail architecture-specific optimizations in our prototype implementation, resulting in more efficient code instrumentation. Specifically, TME-Box integrates the following three compiler modifications:

CPU Register. First, we need to reserve a dedicated CPU register to securely store the base address of the sandbox and enable fast access. TME-Box implements two options for this: We either apply an x86-specific optimization by using a segment register² [39] with segment-relative memory accesses or reserve a general-purpose register. Our compiler framework supports both options, allowing for maximal compatibility and performance comparison, which is reflected in our evaluation.

Depending on the chosen TME-Box mode, we either reserve the `gs` segment register or a single general-purpose register, in our case `r15`, to store the sandbox’s base address (*i.e.*, the upper part of the virtual address) that corresponds to the sandbox’s assigned keyID in the PTEs, which maps to the designated encryption key of the sandbox. Note that reserving a general-purpose CPU register has performance implications due to increased register pressure. Future TME-Box implementations can use the Intel advanced performance extensions (APX)³ [37] that enhance x86 processors with an

²The x86 architecture supports segmentation, allowing memory instructions to use a segment-based addressing mode. In particular, the `fs` and `gs` segment registers are still functional in 64-bit mode. While the `fs` segment is commonly used to address thread local storage (TLS), the `gs` segment has no common use and, thus, can be utilized by applications [39].

³Intel Advanced Performance Extensions (APX) [37] enhance x86 processors by increasing the number of general-purpose registers from 16 to 32.

increased register set. This optimization would then also lessen the overhead of reserving a general-purpose register.

Memory Operations. Next, we integrate a compiler pass that transforms memory operations, enabling the control of the base address and index of memory accesses. The compiler instruments every memory operation to enforce that accesses use the sandbox’s designated encryption key (through the sandbox’s keyID). For instructions that operate with memory (e.g., loads and stores), the corresponding memory operands are truncated, and the sandbox’s base address is instrumented. TME-Box is parameterizable, depending on the size of the virtual address space and the machine’s available TME-MK keys. Depending on the number of keys, the pointer is truncated by clearing the uppermost canonical bits of the address, starting with the highest order bit available. Subsequently, the sandbox’s base address, stored in the reserved register (*i.e.*, either `gs` or `r15`), is instrumented. This is achieved through segment-based addressing or instruction insertion.

Our prototype uses 57-bit virtual addressing, and the Intel Xeon Gold 6530 processor used in our evaluation supports 6-bit keyIDs. In this configuration, we clear the topmost 16 bits of a 64-bit pointer, resulting in 48-bit addresses within sandboxes, leaving enough canonical bits to place the base address in the virtual address pointer. Note that TME-MK supports up to 15-bit keyIDs, and the number of implemented encryption keys is platform-dependent. Since the CPU used in our evaluation supports 6-bit keyIDs, we can operate 63 sandboxes (excluding the default keyID 0). However, for future hardware that supports more keyIDs, our compiler framework can be reconfigured to truncate up to 15 bits of the canonical address to encode the sandbox base address. This truncation is performed by applying a mask to the pointer.

Special care is required for the `eflags`⁴ register of the x86 architecture. The compiler pass checks whether the targeted instructions use `eflags` before instrumenting an instruction sequence. If so, we rely on an instruction sequence that does not alter the `eflags`, as saving and restoring them would induce non-negligible performance overheads. For the `gs` mode, LLVM provides an interface to transform memory operations to segment-based addressing. In `r15` mode, the base address is inserted with an additional instruction. We also implement a common optimization for stack pointer relative accesses similar to existing work [73], [89]. To use this optimization, we need to instrument the register `rsp` and `rbp` whenever they are restored from the stack, as they could be tampered with. Besides stack-relative accesses, all other memory operations are instrumented.

Control-Flow Transfers. Lastly, we implement a compiler pass that isolates the code of the sandbox by instrumenting control-flow transfers. Precisely, forward-edge and backward-edge transitions are instrumented via address masking, restricting transfers to remain in the sandbox’s code region, comparable to recent SFI-based sandboxing [73], [89].

⁴The `eflags` register of the x86 architecture contains the current CPU state represented by status flags, control flags, and system flags [35], [36].

For forward-edge transfers (*i.e.*, indirect function calls and jumps), the compiler instruments the corresponding function pointers before executing the dedicated instruction. This restricts addressable call sites to the virtual address space of the respective sandbox. Similarly, backward-edge transfers (*i.e.*, function returns) are instrumented. Here, our compiler transforms returns into a code sequence that receives the return address from the stack, instruments it, and performs an indirect jump. As an optimization, a future TME-Box implementation can use Intel CET⁵ [79] for return address protection, increasing security and enhancing system performance. Additionally, direct function calls are checked on compile-time for valid call sites, *i.e.*, the sandbox is allowed to call this function directly.

B. Memory Allocator

The memory management of TME-Box is responsible for initializing memory locations of runtime data (e.g., stack and heap memory) of individual sandboxes. This initialization is achieved by writing the entire 64B cache line with the associated keyID. For stack and static memory, this is done during the start-up procedure of the sandbox. The sandbox’s stack memory region is mapped and initialized with the corresponding keyID of the sandbox. Moreover, static memory is copied into the sandbox’s dedicated virtual memory region, *i.e.*, the region is mapped and the corresponding keyID is assigned. Also, the sandbox base address is set up in the dedicated register, *i.e.*, either `gs` or `r15`, of the sandbox.

Furthermore, the memory allocator is responsible for initializing and managing heap memory. Therefore, the memory allocator maps the heap for every sandbox (with their designated keyID) and aliases it to the same physical memory region. Moreover, the memory allocator then performs the initialization of the corresponding heap memory on allocation. This procedure enables the memory allocator to apply the scalable isolation and flexible relocation of data in memory. Precisely, the memory allocator can use fine-grained isolation, ranging from individual cache lines to a large number of pages, depending on the runtime constraints. On pages that contain data from multiple sandboxes, the allocator needs to align memory chunks to cache line size. Here, our design enables the relocation of data, as individual memory chunks can be immediately reallocated. Note that data from different sandboxes is never co-located within a single cache line.

C. Linux Kernel Patch

The operating system needs to provide software support to control the TME-MK hardware feature. For our implementation, we use an experimental patch for the Linux kernel provided by Intel Labs that enables the setting of keyIDs in the PTEs [44]. This kernel patch provides a `syscall` interface that allows additional arguments for the `mprotect` system call to associate a specific keyID with a page (akin to Intel MPK’s `pkey_mprotect`).

⁵Intel Control-flow Enforcement Technology (CET) [79] integrates a shadow stack feature that applies return address protection in hardware.




Fig. 6: The page table entries on the Intel x86-64 architecture [36]. Intel TME-MK applies changes to the specification of physical address [32]. In particular, TME-MK repurposes the upper bits of the physical address to carry the keyID to the encryption engine in the memory controller, resulting in a reduction of addressable physical memory.

Figure 6 illustrates the PTEs on the Intel x86-64 architecture [36], including the changes in the specification to support TME-MK [32]. In particular, TME-MK repurposes the upper bits of the physical address, starting with the highest order bit available, to encode the keyID into the PTE. Thereby, the keyID can be transferred to the encryption engine in the memory controller, performing the encryption procedure. Note that this reduces the addressable physical memory by the number of keyID bits in use.

VI. SECURITY ANALYSIS

In this section, we provide an in-depth security analysis of our TME-Box design. We comprehensively analyze potential security threats and detail the derived security properties of the TME-Box sandbox against an attacker, defined in our thread model (cf. Section III).

A. Systematic Analysis

The generic attack path consists of one or several memory safety vulnerabilities present in the target unprivileged user space application. We assume that the exploitable memory safety error grants the attacker arbitrary read and write capabilities. The attacker can then use this primitive in an attempt to leak or corrupt data in memory, exploiting the software system.

In-Process Memory Isolation. We use TME-Box to isolate individual software components of the application to minimize the attack surface of software exploitation. Consequently, the memory safety error cannot be exploited to escape the isolated sandbox, *i.e.*, the attacker cannot leak or corrupt the memory of other sandboxes.

We achieve this by leveraging compiler instrumentation. TME-Box ensures that the sandbox uses its designated encryption key for memory operations. By controlling the base and index of memory operands, accesses use the sandbox’s keyID, which corresponds to the sandbox’s encryption key. Note that no memory operation exists within a sandbox that bypasses this instrumentation, as this would bypass the protection of the sandbox (*i.e.*, a sandbox escape). Furthermore, we detect memory accesses outside the sandbox’s assigned memory with TME-MK’s cryptographic integrity checks since the sandbox is forced to access memory with its encryption key. Specifically, the TME-MK encryption engine verifies the sandbox’s

memory access. Thus, a violation is detected by the MAC authentication procedure.

Moreover, TME-Box allows for flexible memory management. When a sandbox allocates memory, it is initialized with the sandbox encryption key through a memory write with the corresponding keyID. Subsequently, access to this memory location is solely granted with that specific encryption key (through the keyID) associated with that sandbox. TME-Box supports isolation granularities ranging from single cache lines to full pages. The smallest granularity supported by our isolation mechanism is a 64B cache line, as this is the granularity of the TME-MK integrity checks. Similarly, when memory has been freed by one sandbox and is allocated by another, the memory location is (re-) initialized with the keyID of the new sandbox. Consequently, this memory location can then be accessed exclusively by the new sandbox.

Cryptographic Integrity. The TME-MK encryption engine provides integrity protection for DRAM memory. Once a memory location (*i.e.*, cache line) is initialized with a specific keyID, it is protected by the MAC, which preserves integrity at the cache line level. During memory reads, the memory controller requests the DRAM data and its associated MAC, verifying the memory access.

Successfully verified memory accesses can result in the caching of the data, with the corresponding cache line marked with the keyID used for the access. Any attempt to access memory with an incorrect keyID is detected by the MAC verification within cryptographic bounds. Specifically, cache lines are tagged with the physical address, and the keyID is a part of that. Thus, accessing cache lines with a different keyID results in a cache miss and leads to a DRAM access. This DRAM access is then authenticated by the encryption engine, which detects the usage of incorrect keyIDs.

Writing to memory with an incorrect keyID corrupts the associated MAC. Although this will not immediately result in the detection of the violation, this corruption is also not security-critical since no secret information can be extracted. Subsequent memory reads from this corrupted memory location will trigger a MAC authentication, thus detecting the corruption of the sandbox’s memory. In such cases, the trusted runtime handles the exception, preventing the attacker from exploiting the system.

TME-Box leverages Intel TME-MK’s integrity enforcement, which is based on the security of a cryptographic MAC. Particularly, TME-MK uses a 28-bit MAC generated with the cryptographically secure hash algorithm SHA-3 [7]. The MAC verification ensures that any memory access outside the sandbox’s assigned memory is detected with the probability of $1 - 2^{-28}$, resulting in an exception. Note that attackers cannot bypass this cryptographic integrity check, as our scheme enforces the use of the sandbox’s corresponding keyID for every memory operation, preventing the exploitation of maliciously crafted pointers. Although a MAC collision could theoretically occur with a probability of 2^{-28} , the attacker would only receive wrongly decrypted, mangled data, thereby preserving the confidentiality of the data in memory. Hence,

an attacker would not obtain meaningful information on the correct unencrypted data.

Control-Flow Integrity. Software sandboxing techniques need to ensure that the executed program remains in its designated code region. Thus, control-flow transfers must be restricted so that indirect function calls and function returns are limited to code within the sandbox. TME-Box enforces these restrictions by limiting control-flow transfers to the sandbox code region, similar to recent software-based fault isolation (SFI) techniques [81]. Indirect function calls and function returns are instrumented, applying address masking to enforce that the call site is within the sandbox code region. Additionally, direct function calls are analyzed during compile time to verify that all direct call targets are valid.

Moreover, software sandboxing typically also ensures that untrusted sandboxed code within the process has only restricted access to system calls. These system calls are then executed by a trusted runtime, which can be invoked by the sandbox through a defined interface. This includes critical system calls, such as `mprotect` and `mmap`, used for managing memory resources. Another attack path for memory corruption consists of gaining an arbitrary read and write primitive followed by remote code execution (RCE), *e.g.*, by invoking the `exec` system call. This is followed by a privilege escalation attack to compromise the entire software system [47], [49], [91]. To mitigate this security threat, the usage of system calls is only allowed for the trusted runtime, *i.e.*, the sandbox has to invoke system calls through the trusted runtime. Also, the no-execution policy on writable memory (*i.e.*, Write-XOR-Execute) prevents code injection attacks [6]. In addition, orthogonal measures such as syscall filtering [71] can be employed to further restrict access to system calls.

Multithreading. Security measures must be designed to operate effectively with multithreaded software systems. Ensuring thread safety is crucial to prevent concurrency attacks, exploiting time-of-check to time-of-use (TOCTTOU) [86] vulnerabilities, where an attacker can misuse a time window between security checks to bypass security measures.

TME-Box addresses this security threat since the TME-MK encryption engine verifies each memory access for read operations in hardware. Moreover, TME-MK updates the MAC for cache lines whenever data is written to memory. Specifically, when the CPU requests memory from DRAM, the memory controller fetches the corresponding data and its associated MAC. TME-MK then verifies whether the memory access uses the correct keyID and subsequently caches the data if the MAC verification is successful. Cache lines are tagged with their respective keyID, enabling access to cached data using the appropriate keyID.

When performing data relocation, the memory location is initialized with a new keyID, forcing a writeback for the data in memory that also updates the MAC. This procedure ensures that access for sandboxes with revoked keyIDs becomes invalid. As a result, any subsequent attempts to access memory locations with the revoked keyID trigger MAC verification errors. Note that the processor ensures cache and




Fig. 7: The relative performance overhead of TME-Box using the `gs`-mode for the SPEC CPU2017 benchmark suite.

TLB coherency, maintaining consistency for the cached data and keyIDs located in the PTEs across all CPU cores.

VII. EVALUATION

This section discusses the evaluation of TME-Box in terms of system performance. Moreover, we evaluate the memory latency of the TME-MK hardware feature.

Evaluation Setup. All evaluations are performed on an off-the-shelf Intel Xeon Gold 6530 processor, as this model supports Intel TME-MK. Our system features the following specifications: each of the 32 cores is equipped with a 48 kB L1D cache, a 32 kB L1I cache, and a 2 MB L2 cache. The cores share a 160 MB L3 cache, serving as last-level cache (LLC). Furthermore, our system uses 512 GB DDR5-4800 memory with ECC. Unless stated otherwise, the system is configured to enable Intel TME-MK with integrity support.

A. Performance Evaluation

In this section, we conduct the performance evaluation of our TME-Box design. For our evaluation, we use the SPEC CPU2017 [10] benchmark suite and NGINX. All benchmarks are compiled with the `-O3` optimization level of `clang`.

SPEC CPU2017 Results. We benchmark one baseline configuration and two TME-Box configurations to showcase the performance overheads. Note that we use the `ref` input of SPEC CPU2017 for all benchmarks. Specifically, we evaluate both configurations of TME-Box: TME-Box (`gs`) with segment-based addressing and TME-Box (`r15`) reserving a general-purpose register, as detailed in the implementation section (cf. Section V). In addition, for both configurations, we further distinguish between data isolation, and code and data isolation. Our compiler toolchain targets the hardening of C and C++ applications; thus, we exclude all Fortran benchmarks.

Figure 7 and Figure 8 illustrate the relative performance overhead of TME-Box in `gs`-mode and `r15`-mode, respectively, for the SPEC CPU2017 benchmark suite. Notably, the TME-Box (`gs`) configuration outperforms the TME-Box (`r15`) configuration significantly. We find that TME-Box (`gs`) imposes a low geomean overhead of 5.2% compared to TME-Box (`r15`), which imposes a geomean overhead of 13.4% for data isolation.




Fig. 8: The relative performance overhead of TME-Box using the `r15`-mode for the SPEC CPU2017 benchmark suite.




Fig. 9: Throughput of NGINX with TME-Box for requesting different file sizes normalized to the unmodified NGINX.

In addition, TME-Box requires restricting control-flow transfers for its security. For both code and data isolation, the overhead for TME-Box (`gs`) increases to 9.7%, while the overhead of TME-Box (`r15`) increases to 17.7%. Control-flow instrumentation is particularly sensitive to the number of function calls, as these control-flow transfers are instrumented. We find that the majority of the overhead from code isolation stems from the return address instrumentation. Specifically, the enabled control-flow instrumentation imposes the greatest increase of performance overhead for benchmarks that perform a larger relative number of function calls and returns.

NGINX Results. Additionally, we perform an evaluation of the NGINX web server (version 1.26.0) using an experimental setup comparable to prior work [58], [84], [87]. This experiment uses ApacheBench (`ab`) to generate requests to receive files of different sizes from the NGINX web server. We compile NGINX with our TME-Box compiler toolchain and run a single NGINX worker pinned to an isolated CPU core. We use ApacheBench to perform and benchmark 2,000,000 requests from one client with increasing file sizes covering 1 kB, 8 kB, 32 kB, 128 kB, and 256 kB. Figure 9 shows the throughput of NGINX with TME-Box in both modes (`gs` and `r15`) for code and data isolation normalized to the performance of an unmodified NGINX version. Here, the decrease in throughput for NGINX ranges from 8.3% to 1.8% for TME-Box (`gs`) and from 8.7% to 1.9% for TME-Box (`r15`). We confirm the observations of prior work [84], [87] that the measured throughput decreases stronger for smaller files than for larger files, *i.e.*, the overhead declines as the file size increases.

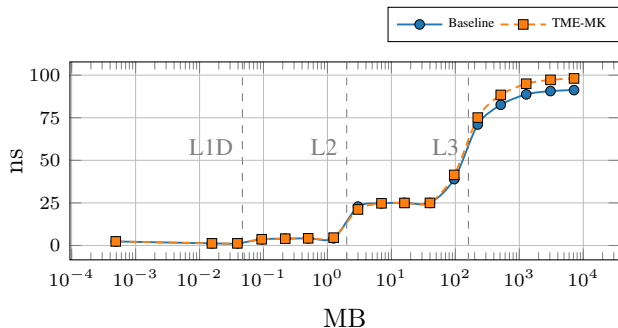


Fig. 10: The memory latency of the Intel TME-MK memory encryption measured with LMBench.

B. Memory Latency

Besides evaluating the performance overhead, we also measure the additional memory latency imposed by the Intel TME-MK memory encryption with integrity. We use the `lat_mem_rd` benchmark of the LMBench [53] suite.

LMBench Results. We perform the benchmark with a memory size of 8 GB and a stride size of 512 B. During benchmarking, the workload is pinned to an isolated CPU core to ensure no other code is executed on that core. Our workload is executed using two configurations. For the first configuration, we enable Intel TME-MK during boot, while the second configuration runs with TME-MK disabled. Figure 10 shows the memory latencies for both variants measured with LMBench. For the L1D, L2, and L3 caches, the access latencies are equivalent. However, as soon as the memory requests are served from DRAM, we observe different latencies for the two configurations. The maximum difference in latency that we observe is 6.8 ns. Note, however, that this overhead is not imposed on each memory access. It only applies to memory operations that cannot be served directly from the cache and must perform a load from DRAM.

VIII. DISCUSSION

This section compares our design with existing work on software isolation and discusses potential extensions.

A. Related Work

In the following, we provide a detailed comparison of TME-Box with prior work, including address space partitioning, memory protection keys, and other isolation mechanisms with page metadata. Additionally, we compare our work with tag-based isolation schemes and approaches that use cryptographic primitives.

Process Isolation. Process isolation techniques, such as memory segmentation and paging, separate the memory resources of different processes that are managed by the operating system. Memory segmentation divides the memory into several regions (*i.e.*, segments for code and data), addressing memory locations using a segment identifier and an offset within the segment. The MMU translates this segment and offset information into a physical address and performs additional

access checks (e.g., read, write, and execute permissions). Thus, segmentation allows the operating system to isolate processes, where access to the memory is only granted if the offset is within the segment length and matching access permissions. Any violation detected by the MMU results in a hardware exception, *i.e.*, a segmentation fault is raised.

Furthermore, paging enables process isolation for modern CPUs through the usage of virtual memory. Paging provides isolation that separates memory access for multiple processes, which is managed on the operating system level. Specifically, virtual addresses are translated by the MMU, which also checks associated access permissions, thereby preventing illegal access to the memory of other processes.

However, the isolation of individual software components with heavy-weight protection mechanisms, such as process isolation, has an impact on system performance and start-up times [3], [66]. Thus, process isolation is not well suited for the isolation of individual software components of applications. As a result, some software systems, like cloud worker architectures [17], require more flexible and efficient in-process isolation that allows scalable memory isolation for different sandboxes, as provided by our mechanism.

Software-based Fault Isolation. Software-based fault isolation (SFI) [42], [85] is a memory isolation technique that divides the virtual address space of a process into predetermined and contiguous regions. Software instrumentation ensures that memory accesses remain within the isolated data and code regions, thereby enforcing in-process isolation for a sandboxed program. Traditional SFI relies on address checking or address masking, while modern SFI approaches use guard zones for this data protection [81].

SFI-based isolation has been implemented and evaluated on modern architectures, such as x86 and ARM [14], [23], [55], [57], [73], [88], [88], [89]. For instance, Google Native Client (NaCl) [73], [89] provides sandboxes of 4 GB in size by leveraging software instrumentation and 40 GB guard zones before and after each sandbox. Moreover, hardware-assisted fault isolation (HFI) [58] proposes a processor extension that implements SFI-style isolation efficiently integrated into the processor’s architecture, thereby reducing the incurred runtime overhead.

While SFI-style approaches allow for a flexible number of sandboxes, they typically cannot provide fine-grained and dynamic isolation, *i.e.*, object-level isolation that can dynamically grow and shrink in size. In addition, SFI is not designed to isolate non-contiguous memory regions. Contrarily, TME-Box allows for scalable and fine-grained (*i.e.*, sub-page granular) isolation of memory with support for flexible data relocation by repurposing the Intel TME-MK encryption engine.

In-Process Isolation with Page Metadata. Protection keys for user space (PKU) is an approach for memory isolation that relies on additional metadata stored in page table entries (PTE). The Intel memory protection keys (MPK) [61] allow to dynamically assign a 4-bit protection key to a specific page. For each memory operation, logical integrity checks are performed that compare the page’s protection key against

the active access policy located and controlled by the user space protection key register (PKRU). The current access permissions (*i.e.*, read and write) of each individual protection key can be changed from user space via the `wrpkru` instruction. In this way, MPK can be used for data protection, *e.g.*, for dynamically locking-away sensitive data in memory like cryptographic key material.

Based on MPK, several academic designs propose software sandboxing [8], [31], [62], [71], [72], [84] that facilitates in-process isolation. However, PKU-based sandboxing can have performance implications in certain scenarios that require frequent sandbox transitions since the `wrpkru` instruction (including necessary fences) can take more than 100 cycles [72], [84]. Similarly, ARMlock [93] enables in-process isolation using the, meanwhile deprecated, ARM memory domains (which provide a 4-bit domain ID for isolation). However, ARM’s policy changes need to be performed by the kernel as the policy register is not accessible from user space, incurring additional runtime overhead for policy changes.

Intel MPK and ARM memory domains only allow for 16 distinct domains, which can be restrictive for certain software architectures that require more in-process domains. Also, both Intel MPK and ARM memory domains are constrained to the granularity of a page and cannot achieve sub-page granular isolation. In contrast to protection keys, TME-Box allows the isolation of more in-process domains and provides fine-grained memory isolation on the level of individual cache lines.

Besides MPK, other academic designs [24], [87] also leverage page metadata located in the PTE for isolation. For example, IMIX [24] proposes a hardware extension that uses a single bit in the PTE to distinguish between secure and non-secure memory pages. Access to data located on secure memory pages is only permitted via newly introduced instructions, thereby preventing memory corruption of security-sensitive pages. CETIS [87] proposes the use of Intel CET [79] in order to separate memory resources by placing the isolated memory region on shadow stack pages that can only be accessed with dedicated instructions (*e.g.*, `wrrss` instruction). These shadow stack pages introduced by Intel CET are identified by an unused combination of the read-write bit and the dirty bit of the PTE. IMIX and CETIS use their isolation primitive to implement metadata protection, such as code-pointer integrity (CPI) [43]. Both countermeasures allow the separation of a process into two distinct domains (*i.e.*, secure and non-secure domains) and cannot provide fine-grained isolation (*i.e.*, sub-page granular protection). While this is enough to protect program metadata efficiently, some software systems require numerous in-process sandboxes (as outlined for MPK).

Tag-based Isolation. Memory tagging technologies associate metadata with data located in memory at a certain granularity, thereby enabling the enforcement of different security policies that restrict access to memory resources [38].

Tagged architectures such as the ARM memory tagging extension (MTE) [75] and SPARC application data integrity (ADI) [4] hardware features have been efficiently integrated into commercial processors. ARM MTE is typically

used to provide probabilistic memory safety for memory sanitization (*e.g.*, MemTagSanitizer) or to help establish runtime security for production software. However, ARM MTE and SPARC ADI are often criticized for their small tag size of 4-bit (*i.e.*, 16 distinct memory tags), resulting in a relatively low detection probability of just 93.75 % [76].

Furthermore, SFITAG [74] combines software instrumentation with ARM MTE to isolate kernel extensions. However, their design only supports a maximum of 14 isolated memory regions, as two memory tags are reserved. In contrast, our TME-Box design supports a larger number of concurrent sandboxes (*i.e.*, TME-MK is specified for up to 15-bit keyIDs that address 32K encryption keys [32]).

Also, HAKC [52] combines ARM pointer authentication (PAuth) [64] with ARM MTE for kernel compartmentalization. While HAKC’s design expands the possible number of compartments from 4-bit MTE, their security is bound to the cryptographic MAC of the PAuth feature, which depends on the size of the virtual address space. More specifically, ARM allows the configuration of virtual address sizes between 32 and 52 bits (with bit 55 reserved). This means that a large virtual address space, as used by servers, can lead to a relatively small MAC, *e.g.*, a 52-bit virtual address space results in a MAC size of 11 bits, or even just 3 bits if ARM’s top-byte ignore (TBI) feature is enabled [64].

Cryptographic Isolation. Cryptographic primitives, such as encryption and authentication, are used in the context of system security to ensure the confidentiality and integrity of code and data located in memory. More specifically, these cryptographic primitives can be applied to mitigate the exploitation of memory safety vulnerabilities [59], [83]. For instance, cryptographic capability computing (C³) [46] leverages pointer and memory encryption to help prevent memory safety errors. C³ encrypts the upper address bits of the pointer, creating a cryptographic address (CA). Furthermore, C³ encrypts and decrypts accessed data in memory using a keystream generator, with the CA serving as input. A memory safety error is detected through C³’s pointer decryption by resolving to a garbled address that likely results in a page fault and subsequent program termination. In contrast to our approach, C³’s proposed hardware modifications directly impact the critical L1 cache latency, whereas TME-Box utilizes Intel TME-MK’s DRAM encryption that is available on commodity x86 CPUs.

Intel TME-MK can be used for runtime security to mitigate various attack vectors [15], [60], [69], [70]. For example, IntegriTag [70] uses Intel TME-MK for probabilistic heap memory safety. IntegriTag uses different security policies (*e.g.*, pseudorandom keyIDs) that select and implicitly encode a keyID into the pointer’s virtual address and assign the keyID to the corresponding memory location through aliasing. This method restricts access to heap objects, as memory accesses are only granted by pointers that incorporate the correct keyID. In this way, TME-MK can be used in the same way as memory tagging, like ARM MTE and SPARC ADI, providing similar security. Nevertheless, adversaries may still exploit vulnerabilities to leak or guess the keyID, and to harvest or forge valid

pointers, thereby circumventing the security measure. In contrast, TME-Box enforces the use of the sandbox’s keyID for all memory operations, enabling scalable in-process isolation. Moreover, EC-CFI [60] uses Intel TME-MK to thwart fault-induced control-flow hijacking attacks. EC-CFI encrypts code at the function level, enforcing decryption with different keys for individual functions. These encryption keys are switched when entering and exiting a function, thereby applying control-flow protection in the presence of fault attacks. EC-CFI helps to prevent fault-induced control-flow hijacking attacks, while our work focuses on software sandboxing.

In addition, Intel trust domain extensions (TDX) [15], [33] enable confidential computing with heavy-weight virtual machine (VM) isolation. TDX uses TME-MK for the encryption of VMs and containers [34], and the protection against physical attacks [30], [33]. In contrast, TME-Box allows for lightweight and efficient in-process sandboxing without relying on the use of heavy-weight virtualization mechanisms.

B. Possible Extensions

Software sandboxing applies coarse-grained CFI to restrict control-flow transfers to the sandbox’s code region. TME-Box adheres to conventional SFI techniques in this regard by instrumenting forward-edge and backward-edge control-flow transitions. Other Intel platform-specific hardware features, *i.e.*, Intel CET’s indirect branch tracking (IBT) and shadow stack [79], enable hardware-enforced CFI. The CET shadow stack feature can be used to ensure the integrity of return addresses. Applying the CET shadow stack for return address protection within the TME-Box framework would be beneficial, thereby increasing security and optimizing performance. In addition, CET’s IBT limits valid indirect jump targets to the landing pads of function entries, thus hardening against control-flow hijacking attacks. IBT could also be useful in the context of TME-Box. However, forward-edge protection still requires instrumentation to enforce that forward-edge transitions remain within the sandbox’s code region. Additionally, FineIBT [25] enforces more fine-grained CFI using IBT. Here, we see potential synergies between TME-Box and fine-grained CFI to further restrict control-flow transfers.

The TME-Box design can be adopted for the isolation of just-in-time (JIT) compiled code since in-process isolation is highly demanded in this context, as seen in the V8 sandbox [27], [29], [68]. To achieve this, the JIT compiler must be TME-Box aware to effectively enable our isolation and memory management, *i.e.*, the JIT compiler must guarantee that the required instrumentation is performed similarly to our compiler extension. Moreover, the JIT compiler has to adhere to the Write-XOR-Execute policy [68], [92], which we assume is enabled for our design. For instance, the Google Native Client (NaCl) [73], [89] sandbox has been extended to the JIT compiler of the V8 engine [5]. We leave it as future work to port TME-Box to JIT compilers.

We use an Intel Xeon Gold 6530 processor as our development and evaluation platform. While TME-MK is specified for up to 15-bit, it is platform-dependent how many are

implemented. For instance, our processor offers support for 6-bit keyIDs. This means that, excluding the default key (keyID 0), TME-Box supports 63 distinct sandboxes on this CPU. Nevertheless, we can still support considerably more sandboxes on such processors by using a hybrid solution of TME-Box and SFI. For instance, in addition to our keyID instrumentation, we could use additional virtual address bits to isolate coarse-grained SFI regions, where the TME-MK keyIDs can be reused for sandboxes in the respective memory region. This allows us to scale our isolation approach up to thousands of sandboxes on processors that do not implement the full 15-bit keyIDs for TME-MK. Thereby, we increase the number of sandboxes while maintaining the advantages of our design, *i.e.*, support for scalable isolation and flexible relocation of data in memory.

IX. CONCLUSION

In this paper, we presented TME-Box, a novel sandboxing technique that provides scalable in-process isolation on commodity x86 CPUs by leveraging Intel’s TME-MK memory encryption. TME-MK is primarily designed to provide page-granular memory encryption for the Intel TDX confidential computing platform, *i.e.*, heavy-weight VM isolation.

TME-Box extends the application of TME-MK’s integrity enforcement to cryptographically isolate the memory of sandboxes *within* a single process (*i.e.*, in-process isolation). That is, TME-Box uses compiler instrumentation to enforce that the sandboxes use their designated encryption keys for memory interactions, thus detecting unauthorized memory accesses. Repurposing TME-MK’s runtime encryption for hardware-assisted sandboxing provides unique isolation properties: TME-Box is scalable, enabling isolation granularities from individual cache lines up to full pages. Moreover, our design allows the isolation of up to 32K concurrent sandboxes and offers flexible memory management for the allocator, *i.e.*, data relocation in memory.

We prototype our TME-Box framework, consisting of an LLVM compiler toolchain and memory allocator, and propose architecture-specific optimizations, *e.g.*, x86 segment-based addressing. Our performance-optimized prototype demonstrates practical results, showcasing geometric performance overheads of 5.2% for data and 9.7% for code and data isolation evaluated using the SPEC CPU2017 benchmark suite.

ACKNOWLEDGMENT

We thank Andreas Kogler and the anonymous reviewers for their valuable feedback that improved this work. This project has received funding from the Austrian Research Promotion Agency (FFG) via the AWARE project (FFG grant number 891092) and the SEIZE project (FFG grant number 888087). Additional funding was provided by a generous gift from Intel. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding parties.

REFERENCES

- [1] Martín Abadi, Mihai Budiu, Úlfar Erlingsson, and Jay Ligatti. Control-Flow Integrity. In *CCS*, 2005.
- [2] National Security Agency. NSA Cybersecurity Information Sheet: Software Memory Safety. https://media.defense.gov/2022/Nov/10/2003112742/-1/-1/0/CSI_SOFTWARE_MEMORY_SAFETY.PDF, 2022. Accessed: 2023-02-26.
- [3] Mark Aiken, Manuel Fähndrich, Chris Hawblitzel, Galen C. Hunt, and James R. Larus. Deconstructing Process Isolation. In *MSPC*, 2006.
- [4] Kathirgamar Aingaran, Sumti Jairath, Georgios K. Konstadinidis, Serena Leung, Paul Loewenstein, Curtis McAllister, Stephen Phillips, Zoran Radovic, Ram Sivaramakrishnan, David Smentek, and Thomas Wicki. M7: Oracle's Next-Generation Sparc Processor. *IEEE Micro*, 35:36–45, 2015.
- [5] Jason Ansel, Petr Marchenko, Úlfar Erlingsson, Elijah Taylor, Brad Chen, Derek L. Schuff, David Sehr, Cliff Biffle, and Bennet Yee. Language-Independent Sandboxing of Just-In-Time Compilation and Self-Modifying Code. In *PLDI*, 2011.
- [6] Iván Arce. The Shellcode Generation. *IEEE Security & Privacy*, 2:72–76, 2004.
- [7] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. The Keccak SHA-3 submission. <http://keccak.noekeon.org/>, 2011. Accessed: 2024-06-10.
- [8] William Blair, William K. Robertson, and Manuel Egele. ThreadLock: Native Principal Isolation Through Memory Protection Keys. In *ASIACCS*, 2023.
- [9] Tyler K. Bletsch, Xuxian Jiang, Vincent W. Freeh, and Zhenkai Liang. Jump-Oriented Programming: A New Class of Code-Reuse Attack. In *ASIACCS*, 2011.
- [10] James Bucek, Klaus-Dieter Lange, and Jóakim von Kistowski. SPEC CPU2017: Next-Generation Compute Benchmark. In *ICPE*, 2018.
- [11] Erik Buchanan, Ryan Roemer, Hovav Shacham, and Stefan Savage. When Good Instructions Go Bad: Generalizing Return-Oriented Programming to RISC. In *CCS*, 2008.
- [12] Nathan Burow, Scott A. Carr, Joseph Nash, Per Larsen, Michael Franz, Stefan Brunthaler, and Mathias Payer. Control-Flow Integrity: Precision, Security, and Performance. *ACM Computing Surveys*, 50:16:1–16:33, 2017.
- [13] Nicholas Carlini, Antonio Barresi, Mathias Payer, David A. Wagner, and Thomas R. Gross. Control-Flow Bending: On the Effectiveness of Control-Flow Integrity. In *USENIX Security*, 2015.
- [14] Miguel Castro, Manuel Costa, Jean-Philippe Martin, Marcus Peinado, Periklis Akrkitidis, Austin Donnelly, Paul Barham, and Richard Black. Fast Byte-Granularity Software Fault Isolation. In *SOSP*, 2009.
- [15] Pau-Chen Cheng, Wojciech Ozga, Enrique Valdez, Salman Ahmed, Zhongshu Gu, Hani Jamjoom, Hubertus Franke, and James Bottomley. Intel TDX Demystified: A Top-Down Approach. *ACM Computing Surveys*, 56:238:1–238:33, 2024.
- [16] Cloudflare. Cloudflare Workers. <https://workers.cloudflare.com/>, 2019. Accessed: 2024-06-10.
- [17] Cloudflare. Cloudflare Security Model. <https://developers.cloudflare.com/workers/reference/security-model/>, 2024. Accessed: 2024-06-10.
- [18] Joan Daemen and Vincent Rijmen. The Block Cipher Rijndael. In *CARDIS*, 1998.
- [19] Joan Daemen and Vincent Rijmen. *The Design of Rijndael: AES - The Advanced Encryption Standard*. Information Security and Cryptography, 2002.
- [20] Deno. A Globally Distributed JavaScript VM. <https://deno.com/deploy>, 2021. Accessed: 2024-06-10.
- [21] Zakir Durumeric, Frank Li, James Kasten, Johanna Amann, Jethro Beekman, Mathias Payer, Nicolas Weaver, David Adrian, Vern Paxson, Michael Bailey, and J. Alex Halderman. The Matter of Heartbleed. In *IMC*, 2014.
- [22] Fastly. Serverless Compute Environment. <https://www.fastly.com/products/edge-compute/serverless>, 2021. Accessed: 2024-06-10.
- [23] Bryan Ford and Russ Cox. Vx32: Lightweight User-level Sandboxing on the x86. In *USENIX ATC*, 2008.
- [24] Tommaso Frassetto, Patrick Jauernig, Christopher Liebchen, and Ahmad-Reza Sadeghi. IMIX: In-Process Memory Isolation EXTension. In *USENIX Security*, 2018.
- [25] Alexander J. Gaidis, Joao Moreira, Ke Sun, Alyssa Milburn, Vaggelis Atlidakis, and Vasileios P. Kemerlis. FineIBT: Fine-grain Control-flow Enforcement with Indirect Branch Tracking. In *RAID*, 2023.
- [26] Jonathan Ganz and Sean Peisert. ASLR: How Robust Is the Randomness? In *SecDev*, 2017.
- [27] Victor Gomes. V8 is Faster and Safer than Ever! <https://v8.dev/blog/holiday-season-2023>, 2023. Accessed: 2024-06-10.
- [28] John Graham-Cumming. Incident report on memory leak caused by Cloudflare parser bug. <https://blog.cloudflare.com/incident-report-on-memory-leak-caused-by-cloudflare-parser-bug>, 2017. Accessed: 2024-06-10.
- [29] Samuel Groß. The V8 Sandbox. <https://v8.dev/blog/sandbox>, 2024. Accessed: 2024-06-10.
- [30] J. Alex Halderman, Seth D. Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A. Calandrino, Ariel J. Feldman, Jacob Appelbaum, and Edward W. Felten. Lest We Remember: Cold Boot Attacks on Encryption Keys. In *USENIX Security*, 2008.
- [31] Mohammad Hedayati, Spyridoula Gravani, Ethan Johnson, John Criswell, Michael L. Scott, Kai Shen, and Mike Marty. Hodor: Intra-Process Isolation for High-Throughput Data Plane Libraries. In *USENIX ATC*, 2019.
- [32] Intel. Intel Architecture Memory Encryption Technologies. <https://www.intel.com/content/www/us/en/content-details/679154/intel-architecture-memory-encryption-technologies-specification.html>, 2022. Revision 1.4, Accessed: 2023-01-31.
- [33] Intel. Intel Trust Domain Extensions. <https://cdrdv2-public.intel.com/690419/TDX-Whitepaper-February2022.pdf>, 2022. Accessed: 2024-05-27.
- [34] Intel. Runtime Encryption of Memory with Intel Total Memory Encryption-Multi-Key (Intel TME-MK). <https://www.intel.com/content/www/us/en/developer/articles/news/runtime-encryption-of-memory-with-intel-tme-mk.html>, 2022. Accessed: 2024-05-27.
- [35] Intel. Intel 64 and IA-32 Architectures Software Developer's Manual Volume 2A: Instruction Set Reference, A-L. <https://www.intel.com/content/dam/www/public/us/en/documents/manuals/64-ia-32-architectures-software-developer-vol-2a-manual.pdf>, 2023. Accessed: 2023-02-26.
- [36] Intel. Intel 64 and IA-32 Architectures Software Developer's Manual Volume 3A: System Programming Guide, Part 1. <https://www.intel.com/content/dam/www/public/us/en/documents/manuals/64-ia-32-architectures-software-developer-vol-3a-part-1-manual.pdf>, 2023. Accessed: 2023-02-26.
- [37] Intel. Intel Advanced Performance Extensions (Intel APX) Architecture Specification. <https://www.intel.com/content/www/us/en/content-details/784266/intel-advanced-performance-extensions-intel-apx-architecture-specification.html>, 2023. Accessed: 2024-01-31.
- [38] Samuel Jero, Nathan Burow, Bryan C. Ward, Richard Skowrya, Roger Khazan, Howard E. Shrobe, and Hamed Okhravi. TAG: Tagged Architecture Guide. *ACM Computing Surveys*, 55:124:1–124:34, 2023.
- [39] The kernel development community. Using FS and GS segments in user space applications. https://www.kernel.org/doc/html/next/x86/x86_64/fsgs.html, 2024. Accessed: 2024-06-10.
- [40] Yoongu Kim, Ross Daly, Jeremie S. Kim, Chris Fallin, Ji-Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu. Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors. In *ISCA*, 2014.
- [41] Paul Kocher, Jann Horn, Anders Fogh, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom. Spectre Attacks: Exploiting Speculative Execution. In *S&P*, 2019.
- [42] Matthew Kolosick, Shravan Narayan, Evan Johnson, Conrad Watt, Michael LeMay, Deepak Garg, Ranjit Jhala, and Deian Stefan. Isolation without Taxation: Near-Zero-Cost Transitions for WebAssembly and SFI. *ACM Programming Languages*, 6:1–30, 2022.
- [43] Volodymyr Kuznetsov, Laszlo Szekeres, Mathias Payer, George Candea, R. Sekar, and Dawn Song. Code-Pointer Integrity. In *OSDI*, 2014.
- [44] Intel Labs. TME-MK-i for Memory Safety. <https://github.com/intellabs/tme-mk-fine-grained-encryption-integrity>, 2024. Accessed: 2024-05-20.
- [45] Chris Latner and Vikram S. Adve. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *CGO*, 2004.
- [46] Michael LeMay, Joydeep Rakshit, Sergej Deutsch, David M. Durham, Santosh Ghosh, Anant Nori, Jayesh Gaur, Andrew Weiler, Salmin Sultana, Karanvir Grewal, and Sreenivas Subramoney. Cryptographic Capability Computing. In *MICRO*, 2021.
- [47] Zhenpeng Lin, Yuhang Wu, and Xinyu Xing. DirtyCred: Escalating Privilege in Linux Kernel. In *CCS*, 2022.

- [48] Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Anders Fogh, Jann Horn, Stefan Mangard, Paul Kocher, Daniel Genkin, Yuval Yarom, and Mike Hamburg. Meltdown: Reading Kernel Memory from User Space. In *USENIX Security*, 2018.
- [49] Lukas Maar, Stefan Gast, Martin Unterguggenberger, Mathias Oberhuber, and Stefan Mangard. SLUBStick: Arbitrary Memory Writes through Practical Software Cross-Cache Attacks within the Linux Kernel. In *USENIX Security*, 2024.
- [50] Luther Martin. XTS: A Mode of AES for Encrypting Hard Disks. *IEEE Security & Privacy*, 8:68–69, 2010.
- [51] Stephen McCamant and Greg Morrisett. Evaluating SFI for a CISC Architecture. In *USENIX Security*, 2006.
- [52] Derrick Paul McKee, Yianni Giannaris, Carolina Ortega, Howard E. Shrobe, Mathias Payer, Hamed Okhravi, and Nathan Burow. Preventing Kernel Hacks with HAKCs. In *NDSS*, 2022.
- [53] Larry W. McVoy and Carl Staelin. Imbench: Portable Tools for Performance Analysis. In *USENIX ATC*, 1996.
- [54] Matt Miller. Trends, challenges, and strategic shifts in the software vulnerability mitigation landscape. https://github.com/Microsoft/MSRC-Security-Research/blob/master/presentations/2019_02_BlueHatIL/2019_01%20-%20BlueHatIL%20-%20Trends%2C%20challenge%2C%20and%20shifts%20in%20software%20vulnerability%20mitigation.pdf, 2019. Accessed: 2023-02-26.
- [55] Greg Morrisett, Gang Tan, Joseph Tassarotti, Jean-Baptiste Tristan, and Edward Gan. RockSalt: Better, Faster, Stronger SFI for the x86. In *PLDI*, 2012.
- [56] Kit Murdock, David F. Oswald, Flavio D. Garcia, Jo Van Bulck, Daniel Gruss, and Frank Piessens. Plundervolt: Software-based Fault Injection Attacks against Intel SGX. In *S&P*, 2020.
- [57] Shравan Narayan, Tal Garfinkel, Evan Johnson, David Thien, Joey Rudek, Michael LeMay, Anjo Vahldiek-Oberwagner, Dean Tullsen, and Deian Stefan. Segue & ColorGuard: Optimizing SFI Performance and Scalability on Modern x86. In *PLAS*, 2022.
- [58] Shравan Narayan, Tal Garfinkel, Mohammadkazem Taram, Joey Rudek, Daniel Moghimi, Evan Johnson, Chris Fallin, Anjo Vahldiek-Oberwagner, Michael LeMay, Ravi Sahita, Dean M. Tullsen, and Deian Stefan. Going beyond the Limits of SFI: Flexible and Secure Hardware-Assisted In-Process Isolation with HFI. In *ASPLOS*, 2023.
- [59] Pascal Nasahl, Robert Schilling, Mario Werner, Jan Hoogerbrugge, Marcel Medwed, and Stefan Mangard. CrypTag: Thwarting Physical and Logical Memory Vulnerabilities using Cryptographically Colored Memory. In *ASIACCS*, 2021.
- [60] Pascal Nasahl, Salmin Sultana, Hans Liljestrand, Karanvir Grewal, Michael LeMay, David M. Durham, David Schrammel, and Stefan Mangard. EC-CFI: Control-Flow Integrity via Code Encryption Countering Fault Attacks. In *HOST*, 2023.
- [61] Soyeon Park, Sangho Lee, and Taesoo Kim. Memory Protection Keys: Facts, Key Extension Perspectives, and Discussions. *IEEE Security & Privacy*, 21:8–15, 2023.
- [62] Soyeon Park, Sangho Lee, Wen Xu, Hyungon Moon, and Taesoo Kim. libmpk: Software Abstraction for Intel Memory Protection Keys (Intel MPK). In *USENIX ATC*, 2019.
- [63] Matthew Prince. Quantifying the Impact of "Cloudbleed". <https://blog.cloudflare.com/quantifying-the-impact-of-cloudbleed>, 2017. Accessed: 2024-06-10.
- [64] Qualcomm. Pointer Authentication on ARMv8.3. <https://www.qualcomm.com/content/dam/qcomm-martech/dm-assets/documents/pointer-auth-v7.pdf>, 2017. Accessed: 2023-02-26.
- [65] Alex Rebert and Christoph Kern. Secure by Design: Google's Perspective on Memory Safety. Technical report, Google Security Engineering, 2024.
- [66] Charles Reis, Alexander Moshchuk, and Nasko Oskov. Site Isolation: Process Separation for Web Sites within the Browser. In *USENIX Security*, 2019.
- [67] Phillip Rogaway, Mihir Bellare, John Black, and Ted Krovetz. OCB: A Block-Cipher Mode of Operation for Efficient Authenticated Encryption. In *CCS*, 2001.
- [68] Stephen Röttger. Control-flow Integrity in V8. <https://v8.dev/blog/control-flow-integrity>, 2023. Accessed: 2024-06-10.
- [69] David Schrammel, Salmin Sultana, Karanvir Grewal, Michael LeMay, David M. Durham, Martin Unterguggenberger, Pascal Nasahl, and Stefan Mangard. MEMES: Memory Encryption-Based Memory Safety on Commodity Hardware. In *SECRYPT*, 2023.
- [70] David Schrammel, Martin Unterguggenberger, Lukas Lamster, Salmin Sultana, Karanvir Grewal, Michael LeMay, David M. Durham, and Stefan Mangard. Memory Tagging using Cryptographic Integrity on Commodity x86 CPUs. In *EuroS&P*, 2024.
- [71] David Schrammel, Samuel Weiser, Richard Sadek, and Stefan Mangard. Jenny: Securing Syscalls for PKU-based Memory Isolation Systems. In *USENIX Security*, 2022.
- [72] David Schrammel, Samuel Weiser, Stefan Steinegger, Martin Schwarzl, Michael Schwarz, Stefan Mangard, and Daniel Gruss. Donky: Domain Keys - Efficient In-Process Isolation for RISC-V and x86. In *USENIX Security*, 2020.
- [73] David Sehr, Robert Muth, Cliff Biffle, Victor Khimenko, Egor Pasko, Karl Schimpf, Bennet Yee, and Brad Chen. Adapting Software Fault Isolation to Contemporary CPU Architectures. In *USENIX Security*, 2010.
- [74] Jiwon Seo, Junseung You, Yungi Cho, Yeongpil Cho, Donghyun Kwon, and Yunheung Paek. Sfitag: Efficient Software Fault Isolation with Memory Tagging for ARM Kernel Extensions. In *ASIACCS*, 2023.
- [75] Kostya Serebryany. ARM Memory Tagging Extension and How It Improves C/C++ Memory Safety. *login Usenix Magazine*, 2019.
- [76] Kostya Serebryany, Evgenii Stepanov, Aleksey Shlyapnikov, Vlad Tsyrvlevich, and Dmitriy Vyukov. Memory Tagging and how it improves C/C++ memory safety. Technical report, Google Security Engineering, 2018.
- [77] Hovav Shacham. The Geometry of Innocent Flesh on the Bone: Return-into-libc without Function Calls (on the x86). In *CCS*, 2007.
- [78] Hovav Shacham, Matthew Page, Ben Pfaff, Eu-Jin Goh, Nagendra Modadugu, and Dan Boneh. On the Effectiveness of Address-Space Randomization. In *CCS*, 2004.
- [79] Vedvyas Shanbhogue, Deepak Gupta, and Ravi Sahita. Security Analysis of Processor Instruction Set Architecture for Enforcing Control-Flow Integrity. In *HASP*, 2019.
- [80] Laszlo Szekeres, Mathias Payer, Tao Wei, and Dawn Song. SoK: Eternal War in Memory. In *S&P*, 2013.
- [81] Gang Tan. Principles and Implementation Techniques of Software-Based Fault Isolation. *Foundations and Trends in Privacy and Security*, 1(3):137–198, 2017.
- [82] Adrian Taylor, Andrew Whalley, Dana Jansens, and Nasko Oskov. An update on Memory Safety in Chrome. <https://security.googleblog.com/2021/09/an-update-on-memory-safety-in-chrome.html>, 2021. Accessed: 2023-02-26.
- [83] Martin Unterguggenberger, David Schrammel, Lukas Lamster, Pascal Nasahl, and Stefan Mangard. Cryptographically Enforced Memory Safety. In *CCS*, 2023.
- [84] Anjo Vahldiek-Oberwagner, Eslam Elnikety, Nuno O. Duarte, Michael Sammler, Peter Druschel, and Deepak Garg. ERIM: Secure, Efficient In-process Isolation with Protection Keys (MPK). In *USENIX Security*, 2019.
- [85] Robert Wahbe, Steven Lucco, Thomas E. Anderson, and Susan L. Graham. Efficient Software-Based Fault Isolation. In *SOSP*, 1993.
- [86] Jinpeng Wei and Calton Pu. TOCTTOU Vulnerabilities in UNIX-Style File Systems: An Anatomical Study. In *FAST*, 2005.
- [87] Mengyao Xie, Chenggang Wu, Yinqian Zhang, Jiali Xu, Yuanming Lai, Yan Kang, Wei Wang, and Zhe Wang. CETIS: Retrofitting Intel CET for Generic and Efficient Intra-process Memory Isolation. In *CCS*, 2022.
- [88] Zachary Yedidia. Lightweight Fault Isolation: Practical, Efficient, and Secure Software Sandboxing. In *ASPLOS*, 2024.
- [89] Bennet Yee, David Sehr, Gregory Dardyk, J. Bradley Chen, Robert Muth, Tavis Ormandy, Shiki Okasaka, Neha Narula, and Nicholas Fullagar. Native Client: A Sandbox for Portable, Untrusted x86 Native Code. In *S&P*, 2009.
- [90] Bin Zeng, Gang Tan, and Greg Morrisett. Combining Control-Flow Integrity and Static Analysis for Efficient and Validated Data Sandboxing. In *CCS*, 2011.
- [91] Kyle Zeng, Zhenpeng Lin, Kangjie Lu, Xinyu Xing, Ruoyu Wang, Adam Doupe, Yan Shoshitaishvili, and Tiffany Bao. RetSpill: Igniting User-Controlled Data to Burn Away Linux Kernel Protections. In *CCS*, 2023.
- [92] Chao Zhang, Mehrdad Niknami, Kevin Zhijie Chen, Chengyu Song, Zhaofeng Chen, and Dawn Song. JITScope: Protecting Web Users from Control-Flow Hijacking Attacks. In *INFOCOM*, 2015.
- [93] Yajin Zhou, Xiaoguang Wang, Yue Chen, and Zhi Wang. ARMlock: Hardware-based Fault Isolation for ARM. In *CCS*, 2014.