

SEV-SNP-strengthening-vm-isolation-with-integrity-protection-and-more-1e2d0f1f63c9168a575bd487c8c8698b

Secure Encrypted Virtualization (SEV)

- minimize their exposure to bugs in the cloud provider's infrastructure
- keep their data confidential

isolate VMs at a hardware level

- individual VMs could be assigned a unique AES encryption key

SEV-ES (Encrypted State)

- protection for CPU register state

SEV-SNP (Secure Nested Paging)

memory integrity protection

- data replay
- corruption
- memory re-mapping
- aliasing based attacks
- if a VM is able to read a private (encrypted) page of memory, it must always read the value it last wrote
- stronger protection around interrupt behavior
- increased protection against recently disclosed side channel attacks
- requires a combination of new CPU hardware and firmware

Threat Model

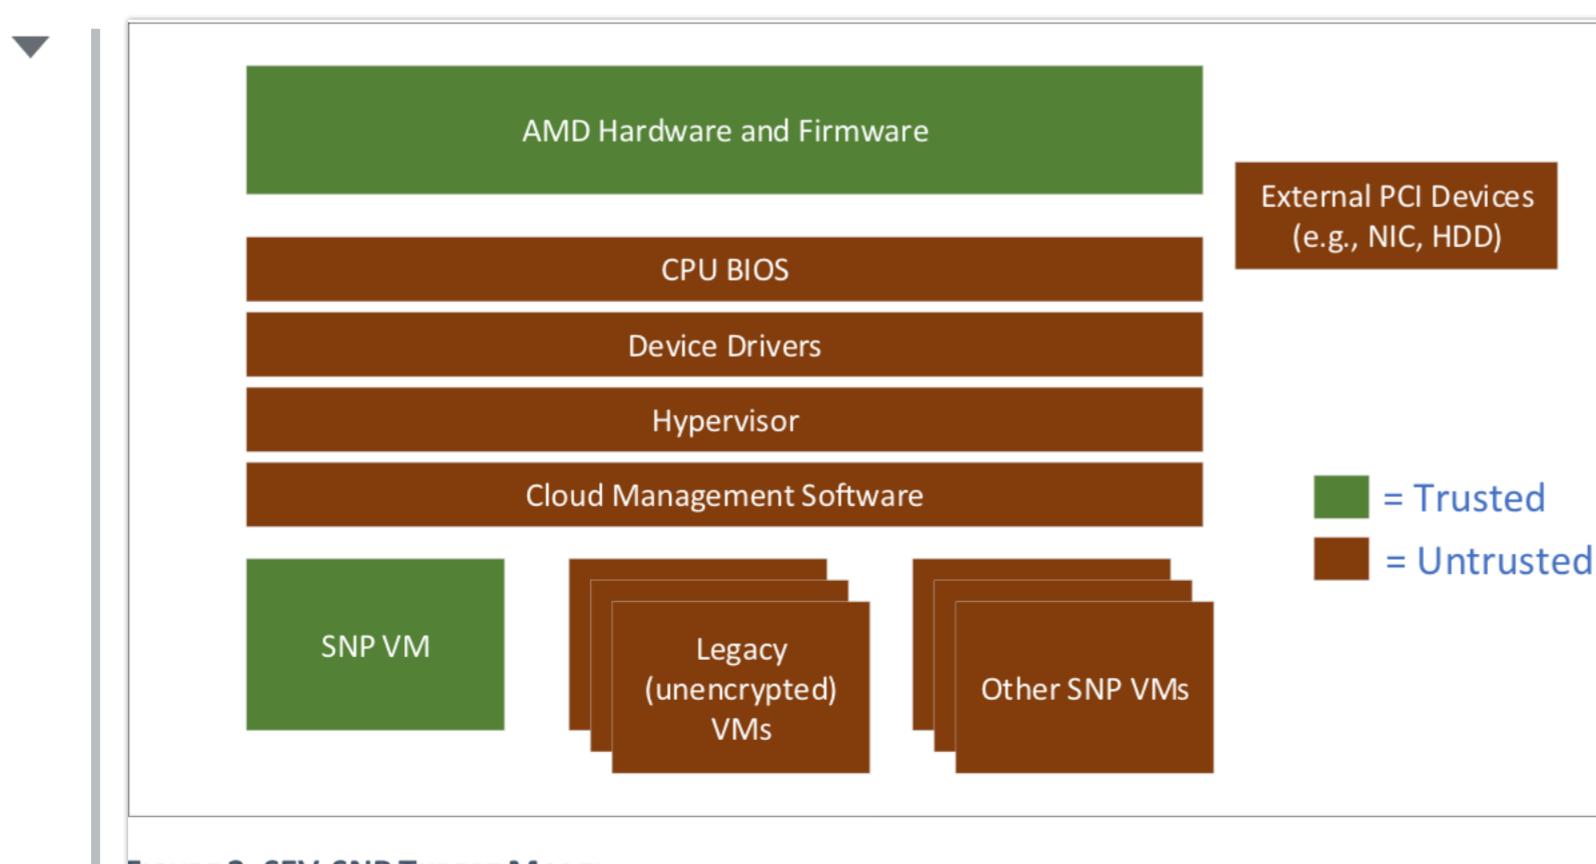


FIGURE 2: SEV-SNP THREAT MODEL

▼ | fully trusted

- | AMD Secure Processor (AMD-SP)

5

- | VM

5

5

▼ | fully untrusted

▼ | all other CPU software components

5

- | BIOS on the host system

5

- | hypervisor

5

- | device drivers

5

- | other VMs

5

- | PCI devices

5

- | SEV technologies only protect data in-use

5

- |

✓ = Mitigated ★ = Optionally Mitigated Ø = Not Mitigated

Potential Threats

Confidentiality

VM Memory

Example attack: Hypervisor reads private VM memory

VM Register State

Example attack: Read VM register state after VMEXIT

DMA Protection

Example attack: Device attempts to read VM memory

Integrity

Replay Protection

Example attack: Replace VM memory with an old copy

Data Corruption

Example attack: Replace VM memory with junk data

Memory Aliasing

Example attack: Map two guest pages to same DRAM page

Memory Re-Mapping

Example attack: Switch DRAM page mapped to a guest page

Availability

Denial of Service on Hypervisor

Example attack: Malicious guest refuses to yield/exit

Denial of Service on Guest

Example attack: Malicious hypervisor refuses to run guest

Physical Access Attacks

Offline DRAM analysis

Example attack: Cold boot

Active DRAM corruption

Example attack: Manipulate DDR bus while VM is running

Misc.

TCB Rollback

Example attack: Revert AMD-SP firmware to old version

Malicious Interrupt/Exception Injection

Example attack: Inject interrupt while RFLAGS.IF=0

Indirect Branch Predictor Poisoning

Example attack: Poison BTB from hypervisor

Secure Hardware Debug Registers

Example attack: Change breakpoints during debug

Trusted CPUID Information

Example attack: Hypervisors lies about platform capabilities

Architectural Side Channels

Example attack: PRIME+PROBE to track VM accesses

Page-level Side Channels

Example attack: Track VM access patterns through page tables

Performance Counter Tracking

Example attack: Fingerprint VM apps by performance data

	SEV	SEV-ES	SEV-SNP
VM Memory	✓	✓	✓
VM Register State	Ø	✓	✓
DMA Protection	✓	✓	✓
Replay Protection	Ø	Ø	✓
Data Corruption	Ø	Ø	✓
Memory Aliasing	Ø	Ø	✓
Memory Re-Mapping	Ø	Ø	✓
Denial of Service on Hypervisor	✓	✓	✓
Denial of Service on Guest	Ø	Ø	Ø
Offline DRAM analysis	✓	✓	✓
Active DRAM corruption	Ø	Ø	Ø
TCB Rollback	Ø	Ø	✓
Malicious Interrupt/Exception Injection	Ø	Ø	★
Indirect Branch Predictor Poisoning	Ø	Ø	★
Secure Hardware Debug Registers	Ø	Ø	★
Trusted CPUID Information	Ø	Ø	★
Architectural Side Channels	Ø	Ø	Ø
Page-level Side Channels	Ø	Ø	Ø
Performance Counter Tracking	Ø	Ø	Ø

7

TABLE 1: THREAT MODEL

Reverse Map Table

- single data structure shared across the system
- one entry for every 4k page of DRAM that may be used by VMs
- tracks the owner for each page of memory
- in conjunction with standard x86 page tables

8

8

8

8

9

- New CPU instructions exist to enable manipulation of RMP entries

10

PVALIDATE

- sets the Validated bit

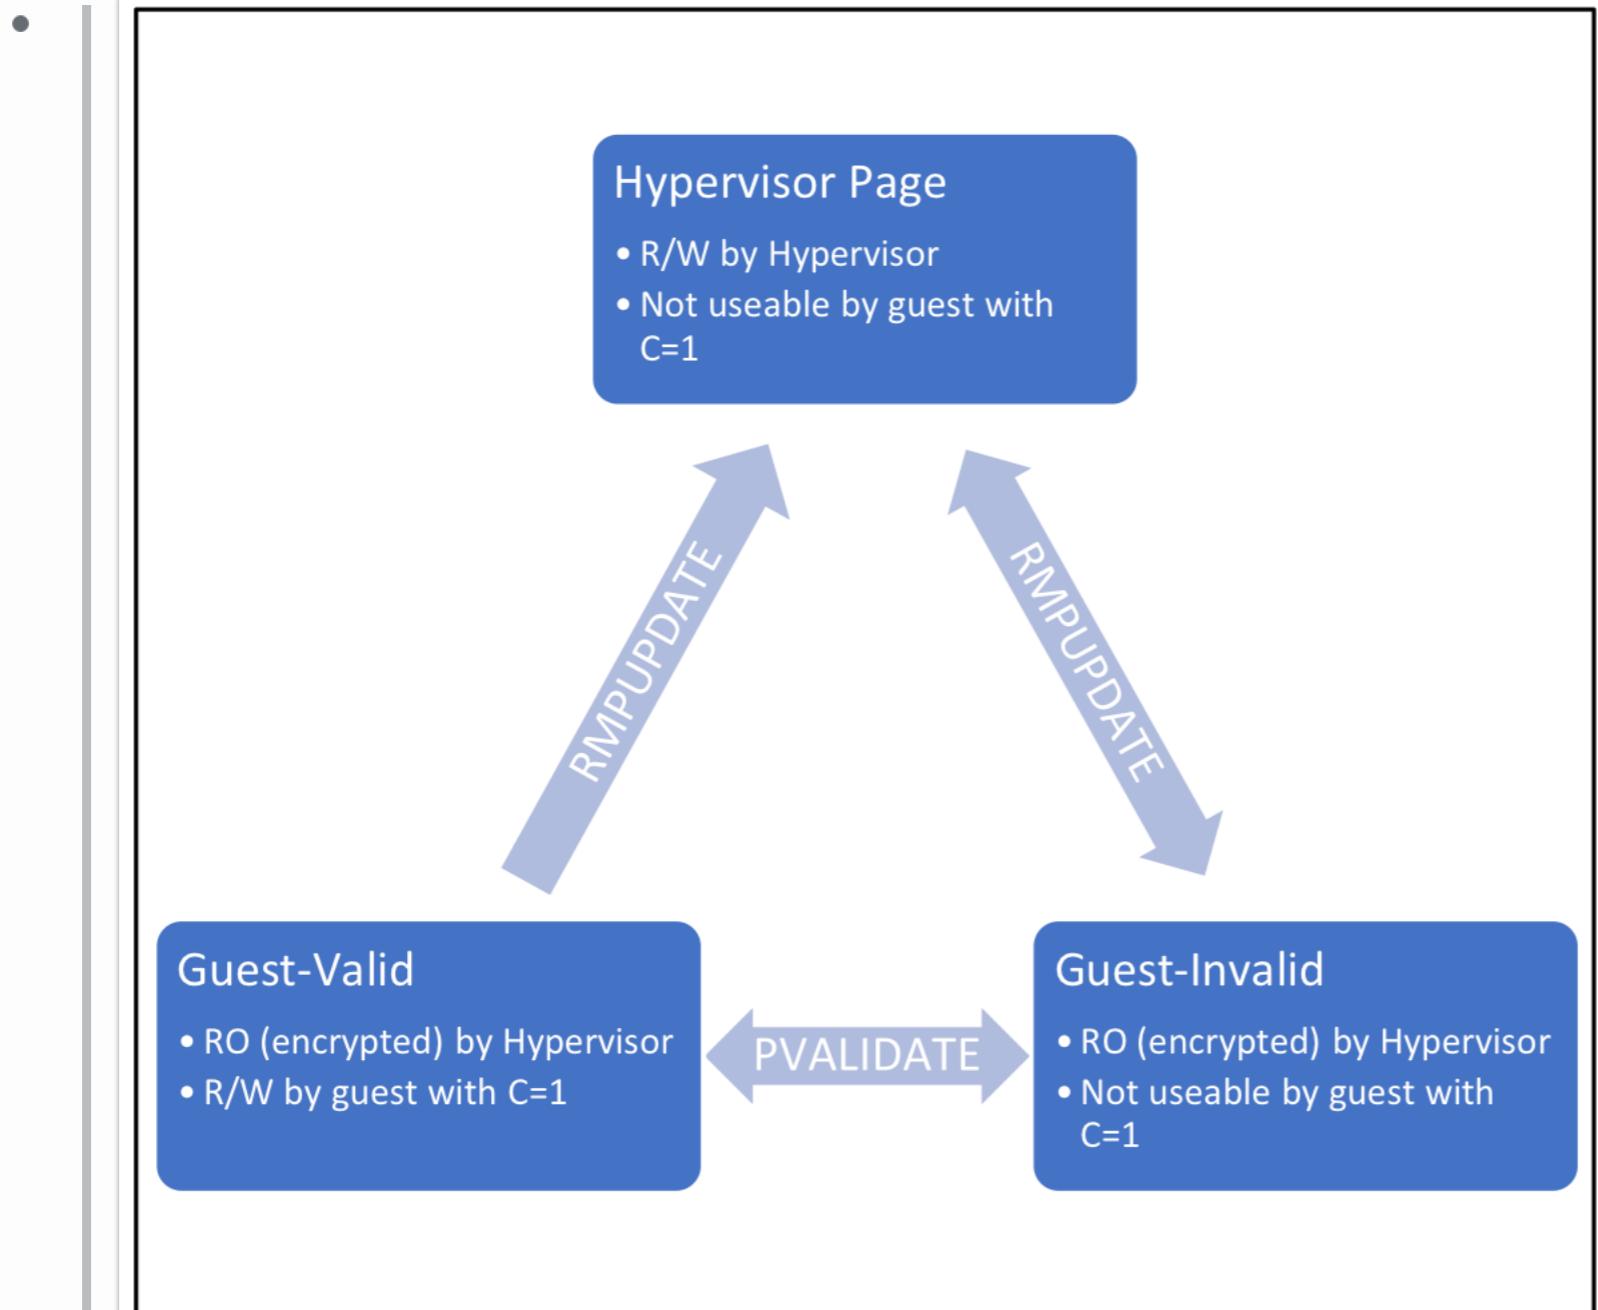
10

10

RMPUPDATE

- assigns the page to the guest

10



10

FIGURE 4: BASIC PAGE STATES

13

RMPADJUST

- allows a given VMPL to modify permissions for a less privileged VMPL
- restricted so one level cannot grant more permissions than it currently has

13

13

AMD-V 2-level paging1

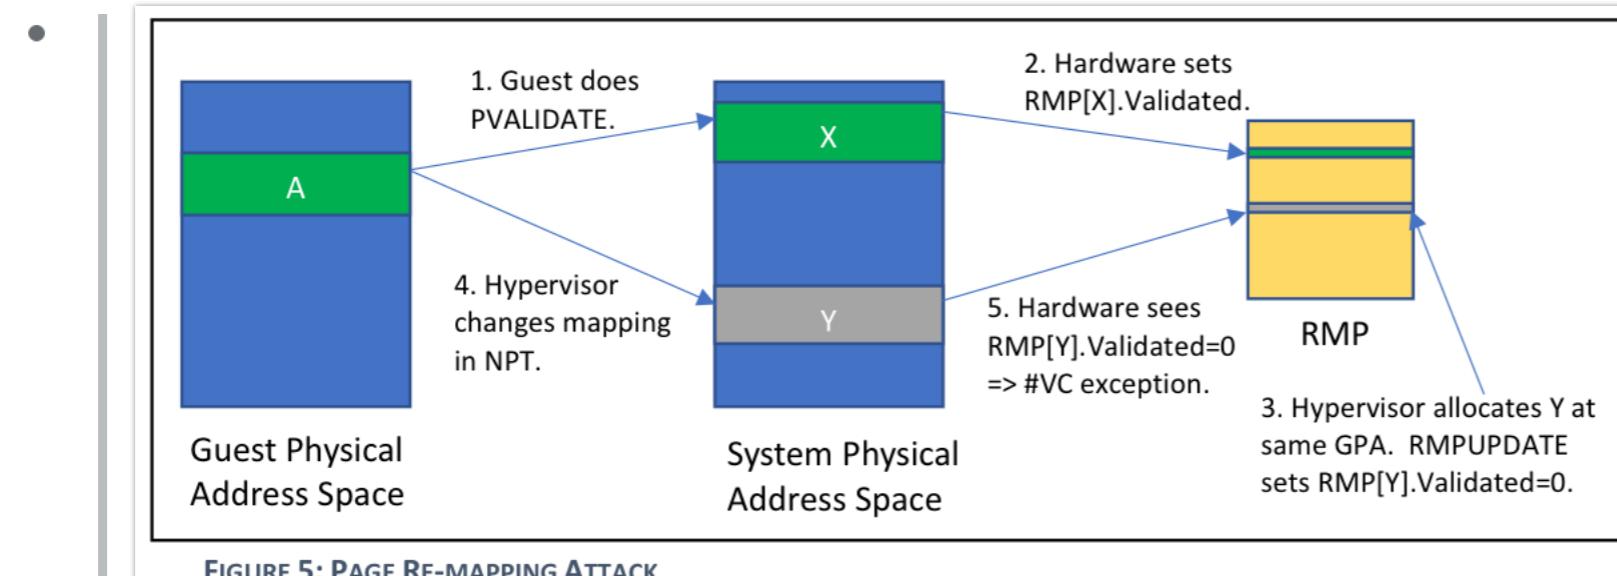
9

translate a Guest Virtual Address (GVA) to a Guest Physical Address (GPA)

- to a System Physical Address (SPA).

- every SPA can only be mapped to a single GPA at one time.

- each RMP entry is a Validated bit



Virtual Machine Privilege Levels

- provide hardware isolated abstraction layers within a VM for additional security controls

- assistance with managing communication with the hypervisor

- RMPADJUST

- allows a guest VM to divide its address space into four levels

- primarily used to set additional page permission checks and are otherwise orthogonal to other x86 security features

VMPL0 is the highest privilege

- handle #VC events

- configure which guest memory in another vCPU is private (C=1) versus shared (C=0)

- intermediary for guest to hypervisor communications

- VMPL3 is the least privileged

Interrupt/Exception Protection

- restrictive interface between the VM and hypervisor regarding interrupts and exceptions

optional modes

- Restricted Injection

- | Alternate Injection 15
- ▼ | Trusted Platform Information 16
 - | Platform features and capabilities 16
 - | malicious hypervisor can only cause denial-of-service on a guest by lying about CPUID features 16
 - | AMD-SP will verify that the CPUID results that the hypervisor is reporting are no greater than the capabilities of the platform and that security sensitive information 16
- ▼ | TCB Versioning 16
 - ▼ | upgradeable firmware components 16
 - | considered trusted 16
 - ▼ | AMD-SP version number 16
 - | SEV-SNP 16
 - | Versioned Chip Endorsement Key (VCEK) 16
 - | unique to each AMD chip 16
- ▼ | VM Launch & Attestation 17
 - | launch digest captures the layout of initial guest memory, as well as its contents 17
 - ▼ | guest owner 17
 - | signed Identity Block (IDB) to associate with the VM 17
 - | uniquely identify the VM 17
 - | can only be associated with VMs which match the expected launch digest 17
 - | included as part of the attestation report 17
 - ▼ | supports more flexible attestation 17
 - | Attestation reports can be requested through a protected path from the AMD-SP by the guest VM at any time 17
- ▼ | VM Migration 18
 - ▼ | Migration Agent (MA) 18
 - | MA is itself an SEV-SNP VM 18
 - | single MA can manage migration for an arbitrary number of VMs 18

- | not required that both the source and destination machines be online at the same time

18

▼ | Side Channels

- | hardware capabilities⁴
- | protection from speculative side channel attacks and SMT
- | not able to protect against all possible side channel attacks
- | PRIME+PROBE
- | it is the responsibility of VM owners to follow standard security practices
- | Modern cryptographic libraries take special care
- | shared (unencrypted) memory

19

19

19

19

19

20

19

4