



AMD64 Enhanced SMT Protection

Publication # **69204**

Revision: **1.00**

Issue Date: **March 2026**

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes. THIS INFORMATION IS PROVIDED "AS IS." AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

AMD, the AMD Arrow logo, AMD EPYC, and combinations thereof are trademarks of Advanced Micro Devices, Inc.

Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies. PCIe[®] is a registered trademark of PCI-SIG Corporation. Linux[®] is the registered trademark of Linus Torvalds in the U.S. and other countries.

© 2026 Advanced Micro Devices, Inc. All rights reserved.

Revision History

| Date | Revision | Change Description |
|------------|----------|--------------------|
| March 2026 | 1.00 | Initial release |

Enhanced SMT Protection

Introduction

The Enhanced SMT Protection (ESMTP) feature extends SMT Protection (SMTP) functionality by allowing the sibling thread(s) to either be in an idle state in host mode or to execute a vCPU from the same VM, which allows more efficient use of hardware multi-threading resources.

Additionally, ESMTP adds the ability for guests to define relationship between vCPUs by specifying which vCPUs may run on the same CPU core.

Presence

CPUID Fn8000_0025_EDX[EnhSmtProtection] (bit 1) indicates ESMTP support.

VMCB and VMSA Changes

The SEV_FEATURES bit 17 (EnhSmtProtection) may be used to enable the ESMTP feature for SNP-active guests. Enhanced SMT Protection and SMT Protection (SEV_FEATURES bit 15) are mutually exclusive. VMRUN fails with VMEXIT_INVALID if both features are enabled.

The VCPU_ID field (VMSA offset 8A0h, 4-bytes) identifies the vCPU. This field should be set during guest creation. An SEV-SNP guest may read the current value of VCPU_ID through the VCPU_ID MSR (C001_013Ah). This MSR is read-only. An attempt to write VCPU_ID MSR will result in a #GP(0) exception. The VCPU_ID MSR read will return zero when ESMTP is not enabled.

The VCPU_SIBLING_MASK field (VMSA offset 8A4h, 4-bytes) defines a mask that determines VCPU_ID values that are allowed to run on the same core simultaneously.

The ESMTP_TIMEOUT_CTL field (VMCB offset 148h, 8-bytes) specifies the wait in P0 clocks to enter guest mode when ESMTP is enabled. ESMTP timeout is disabled if ESMTP_TIMEOUT_CTL is equal to 0.

VMRUN ESMTP Behavior

When a VMRUN is executed to an SEV-SNP (SNP-Active) vCPU with ESMTP enabled, the processor waits until all sibling threads have done a VMRUN to a legal ESMTP sibling (as defined below) or are in an idle state in host mode before entering guest mode. VMRUN to an ESMTP enabled vCPU fails if:

- Any sibling thread does a VMRUN to an illegal ESMTP sibling
- A physical INTR, NMI, SMI, or INIT occurs
- Timeout specified by ESMTP_TIMEOUT_CTL occurs

If a physical interrupt occurs, the VMRUN instruction finishes with the exit code for the physical interrupt (e.g., VMEXIT_INTR). Physical INTR, NMI, SMI, and INIT interrupts detected during VMRUN are unconditionally intercepted. If an internal event is detected during VMRUN, it exits with VMEXIT_RETRY exit code (AE, exit code -7). The hypervisor should re-execute VMRUN if it exits with VMEXIT_RETRY.

If the ESMTP timer expires, the VMRUN instruction exits with VMEXIT_ESMTP_TIMEOUT exit code (AE, exit code -6).

A legal ESMTP sibling is defined as follows:

- SEV-SNP VM with ESMTP enabled
- ASID values match
- VCPU_SIBLING_MASK values match
- VCPU_ID & ~VCPU_SIBLING_MASK values match

An illegal ESMTP sibling is an SEV-SNP vCPU with ESMTP enabled which does not meet all of the above conditions. If any of the above checks fail, VMRUN exits with VMEXIT_ILLSIB exit code (AE, exit code -5).

Once all sibling threads are either in an idle state in host mode (HLT, MWAIT, MWAITX at CPL0, or an I/O C-state) or have executed VMRUN to a legal ESMTP sibling, any pending VMRUN operations on that core are allowed to finish and enter guest mode.

VMEXIT ESMTP Behavior

When an automatic exit event occurs while a vCPU with ESMTP enabled is running, the VMEXIT cannot complete until all sibling threads not in an idle state in host mode stop executing guest code. During the VMEXIT flow, the processor checks the status of the sibling thread(s). If a sibling thread is executing guest code, the processor writes the APIC ICR register with the IDLE_WAKEUP_ICR MSR value and waits until the sibling thread enters host mode. The IPI is not sent if all other sibling threads are in an idle state or are attempting to complete VMEXIT.

If the VMEXIT occurs during the VMRUN instruction, the processor writes the APIC ICR register with the IDLE_WAKEUP_ICR MSR value if any sibling thread with ESMTP enabled is attempting to complete VMRUN or is executing guest code, and waits until the sibling thread enters host mode.