Linux
Plumbers
Conference | Richmond, VA | Nov. 13-15, 2023

Secure TSC for AMD SEV-SNP guests

Nikunj A. Dadhania

Timestamp Counter (TSC)

- A counter implemented in every x86 microprocessor
- Counts processor-clock cycles
- 64-bit register called TSC MSR (0x0010h) provides the latest value of the counter
- RDTSC/RDTSCP instruction can be used to read the Time-Stamp Counter

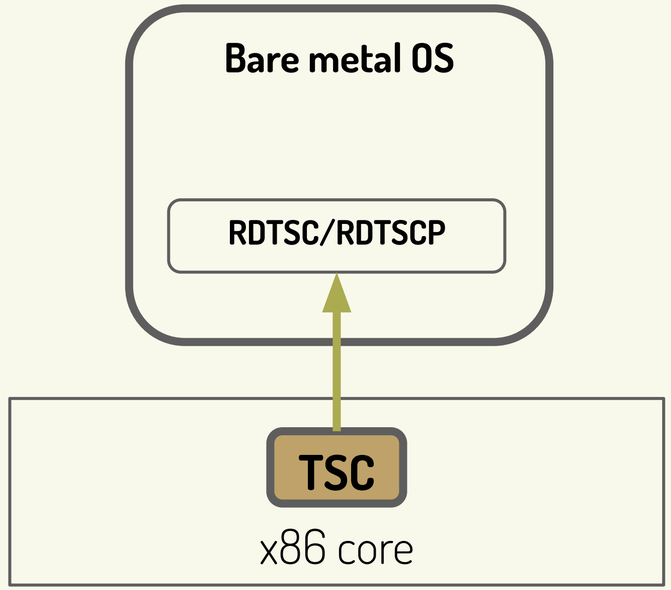

TSC calculation in SVM Guests

- A TSC read in the guest today is computed with following inputs
  - TSC_RAW
  - TSC Ratio MSR C000_0104h
  - VMCB TSC_OFFSET
- Writing to the TSC Ratio MSR and TSC_OFFSET allows the hypervisor to control the guest's view of the TSC.
- **Security concern:** TSC reporting is hypervisor controlled, a malicious hypervisor can prevent guest TSC from moving forward.
- In SEV / SEV- ES guest as well, hypervisor can change the guest's TSC view

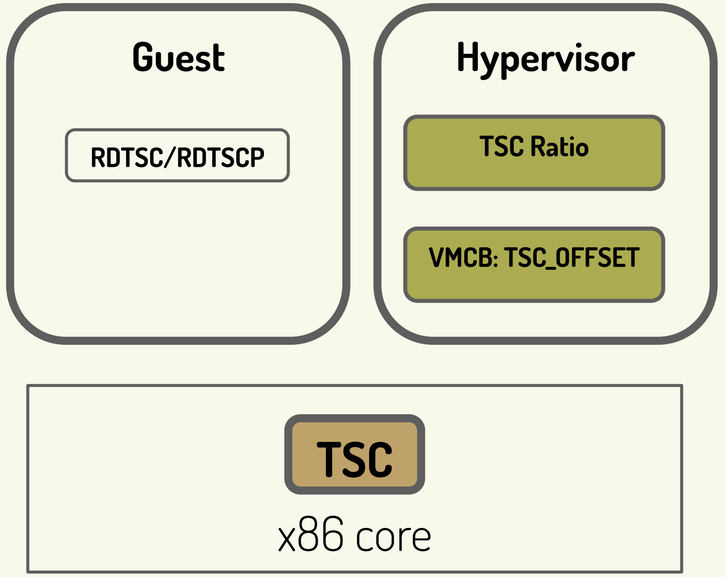

TSC Value (in guest) = (P0 frequency * TSCRatio * t) + VMCB.TSC_OFFSET + (Last Value Written to TSC) * TSCRatio

TSCRatio = (Desired TSCFreq) / Core P0 frequency

Where t is time since the TSC was last written via the TSC MSR (or since reset if not written)

AMD64 Architecture Programmer's Manual - Section "TSC Ratio MSR (C000_0104h)"

SEV-SNP - Secure TSC feature

- Secure TSC enabled SEV-SNP guest does not depend on hypervisor-controlled parameters
- The TSC value in a SecureTSC enabled SNP guest is calculated as follows:

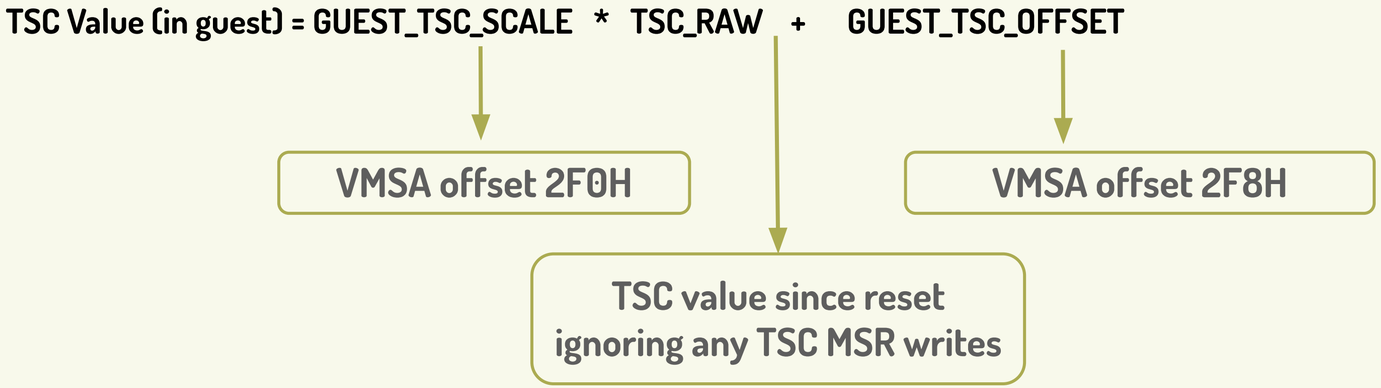

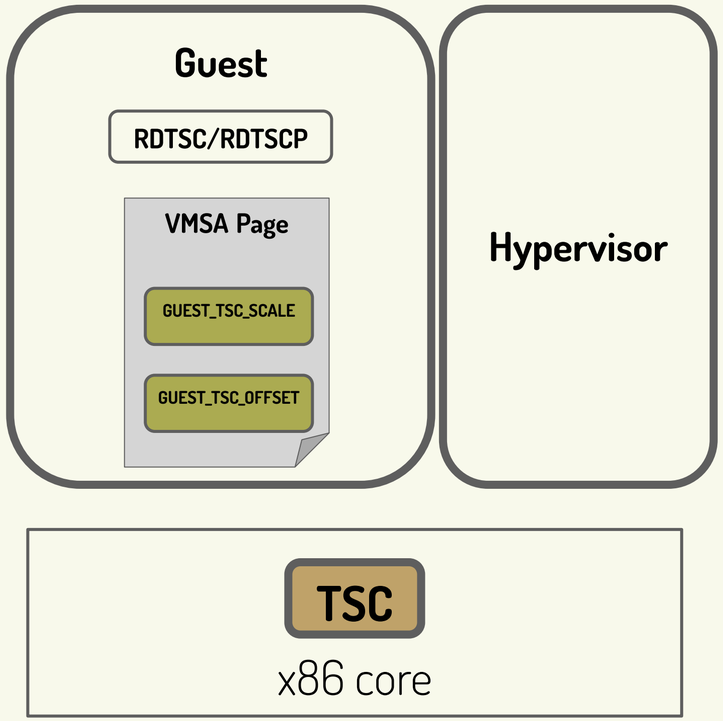

- Parameters are stored in the guest's secure save area - VMSA is an encrypted page.
- GUEST_TSC_FREQ MSR (C001_0134h) read-only MSR provides effective frequency in MHz
- RDTSC/RDTSCP interceptions are prevented by the guest #VC handler
- TSC MSR reads are emulated by the #VC handlers and writes are prevented
- SecureTSC can be used as a clocksource instead of KVM paravirt clock

Boot sequence with SecureTSC

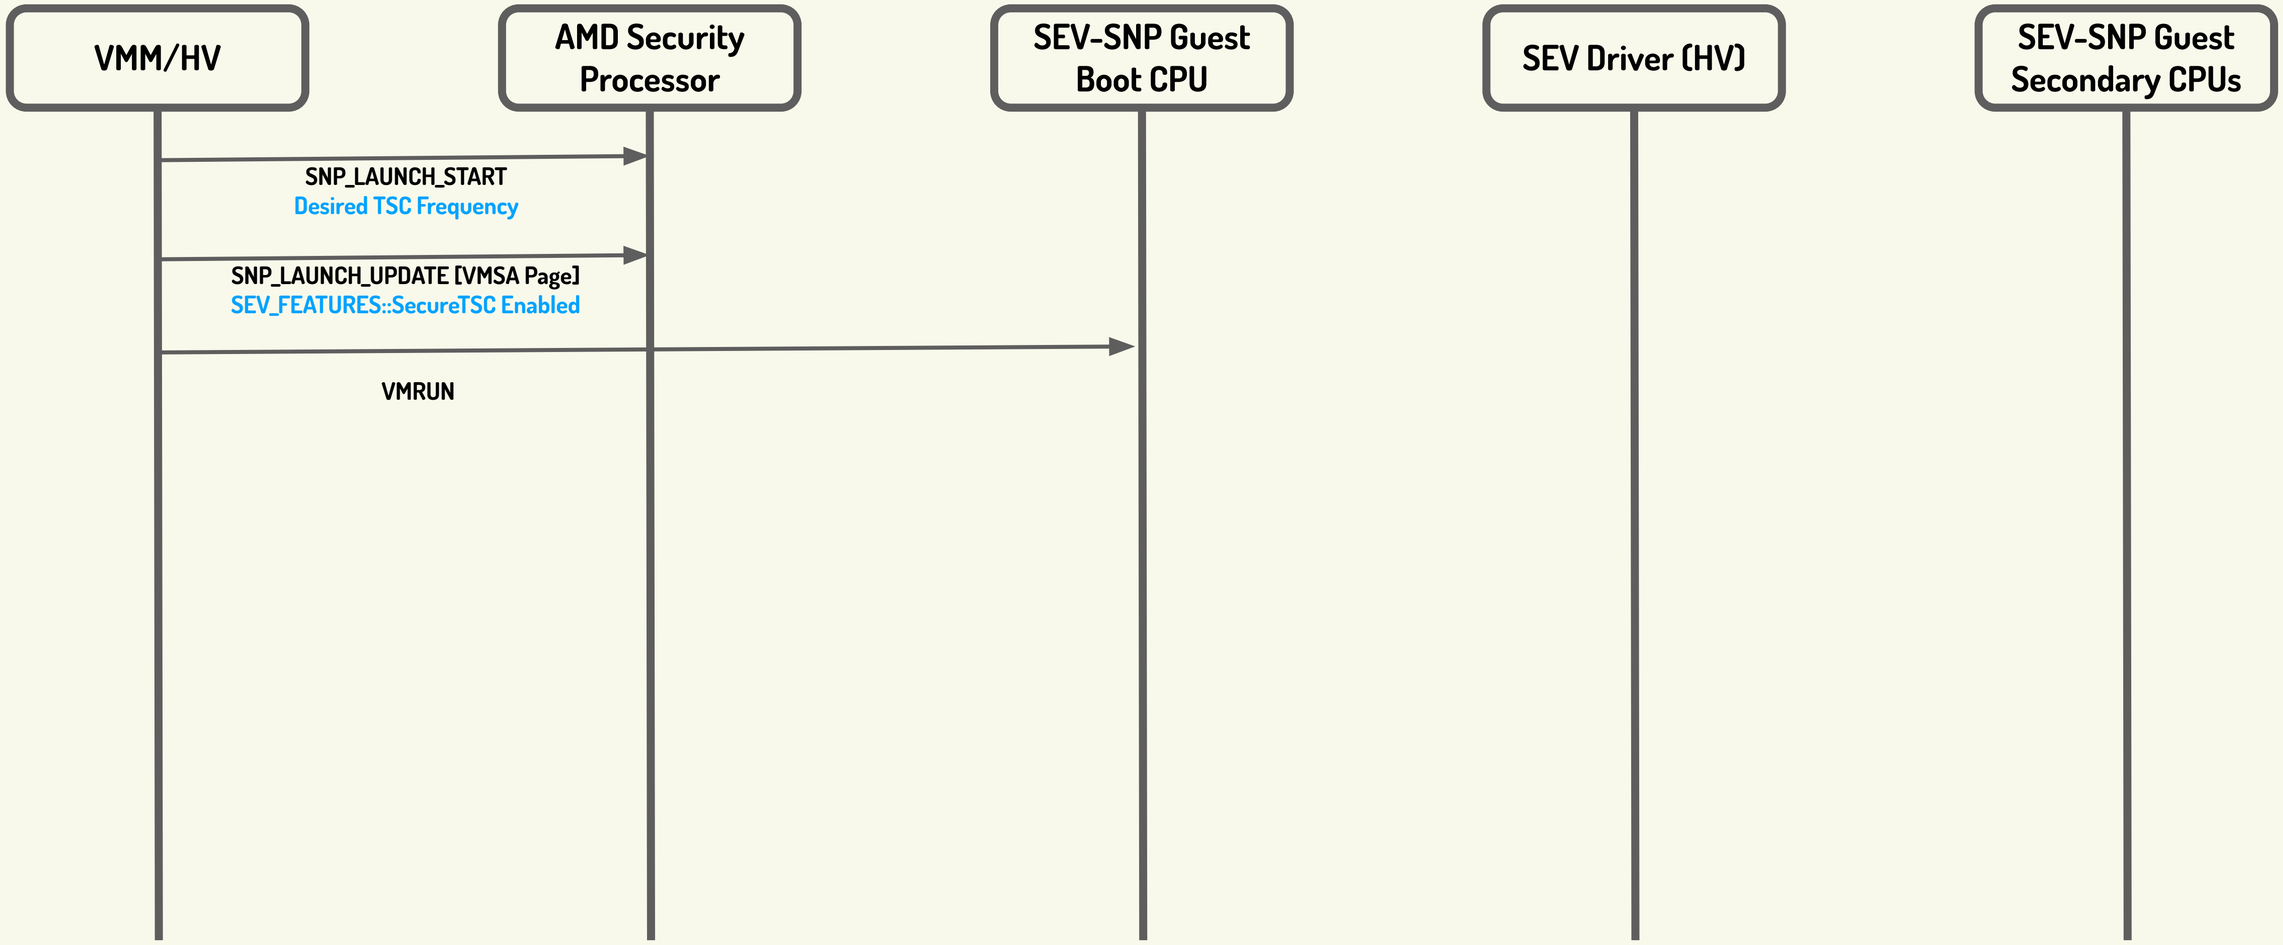

Simplified view

Boot sequence with SecureTSC

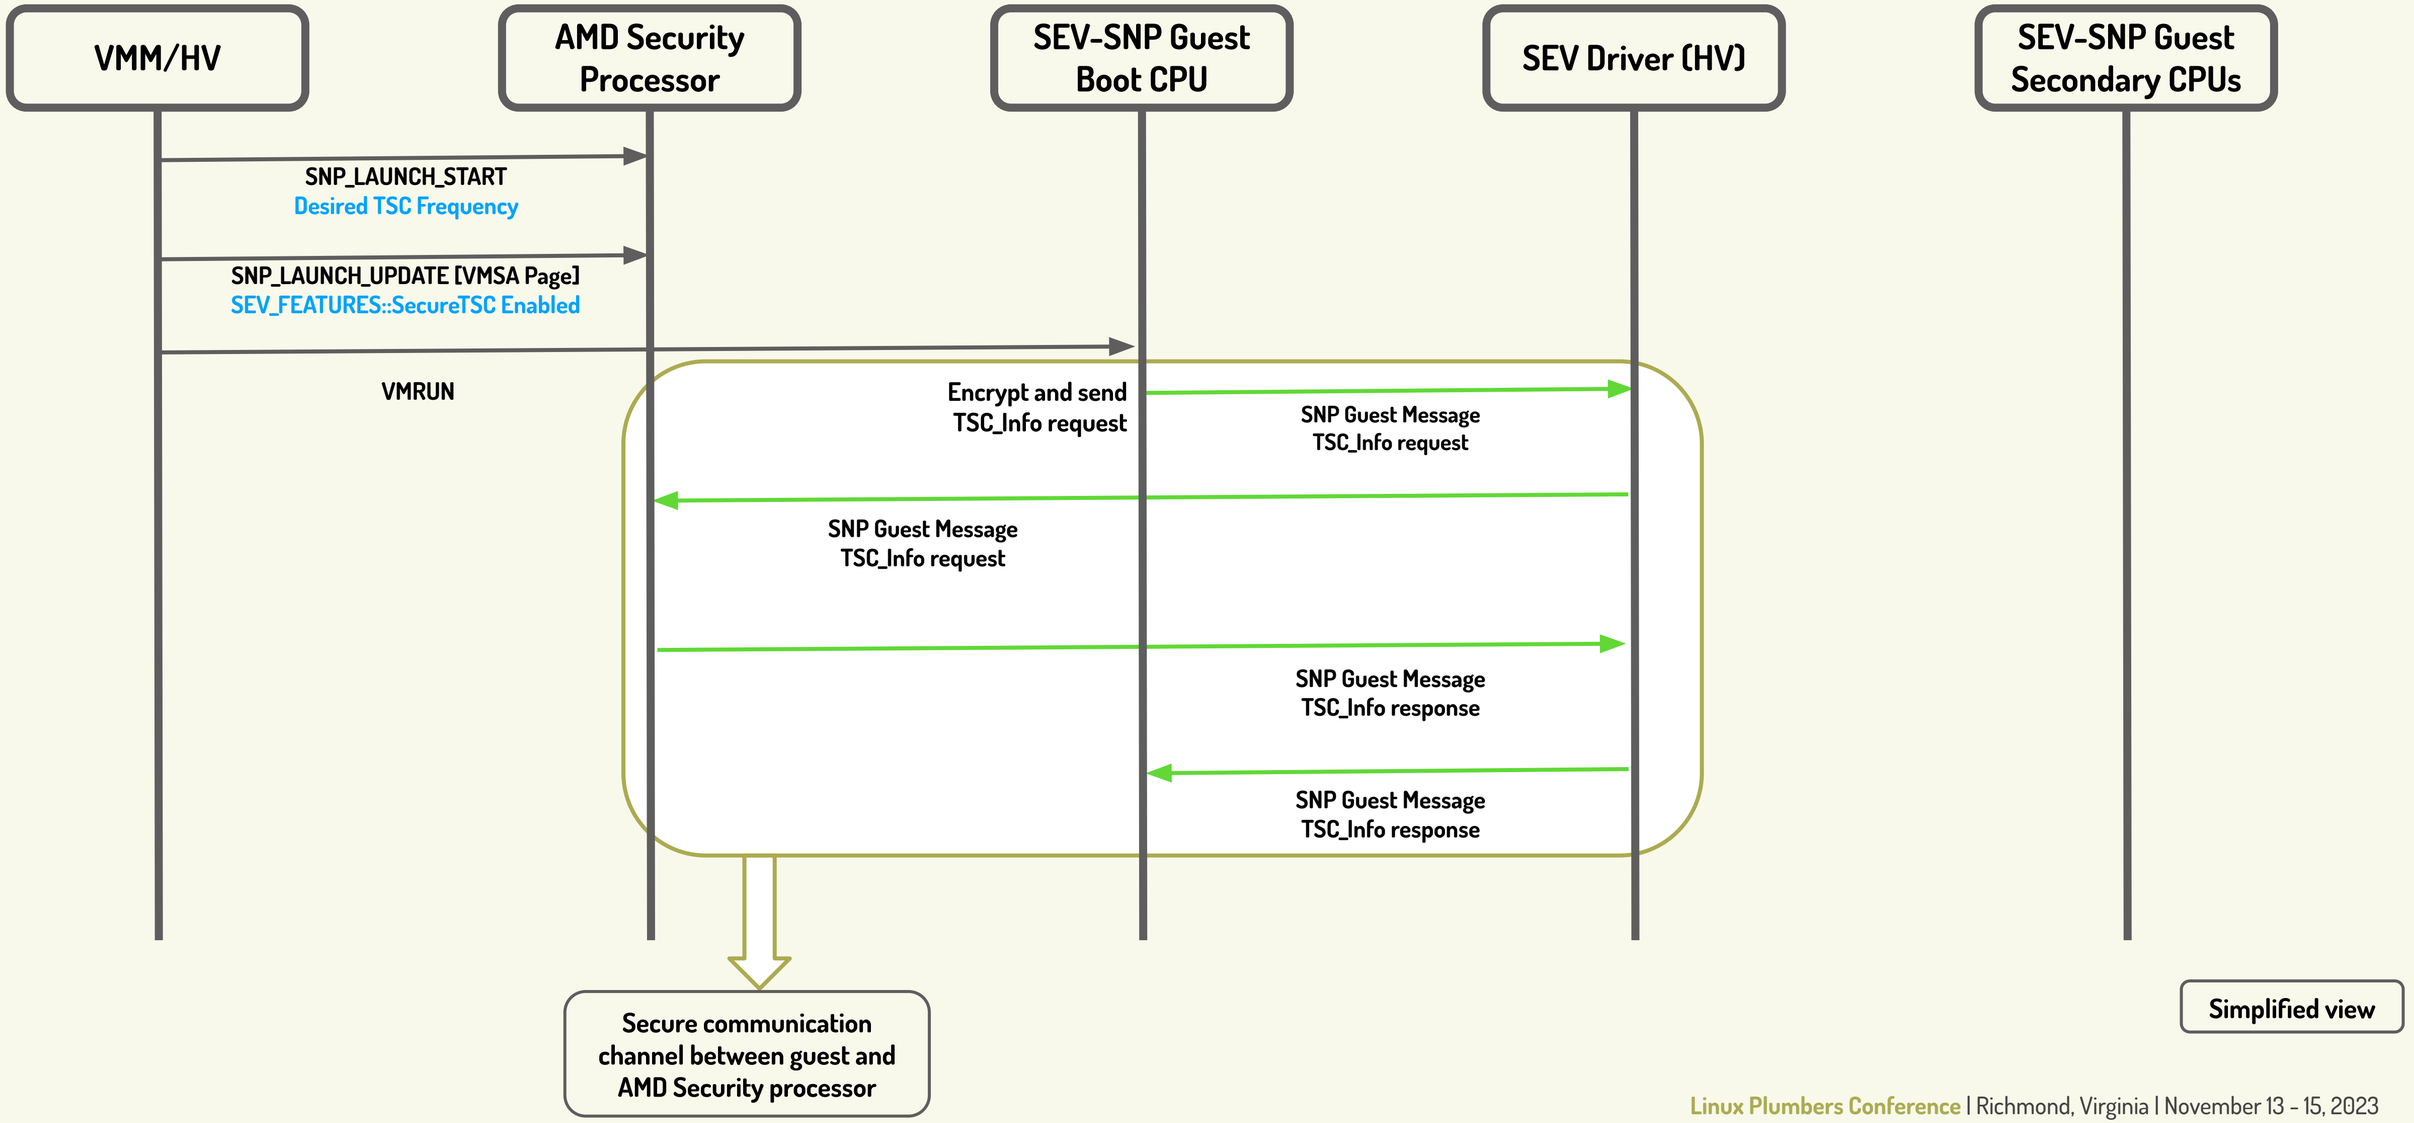

Boot sequence with SecureTSC

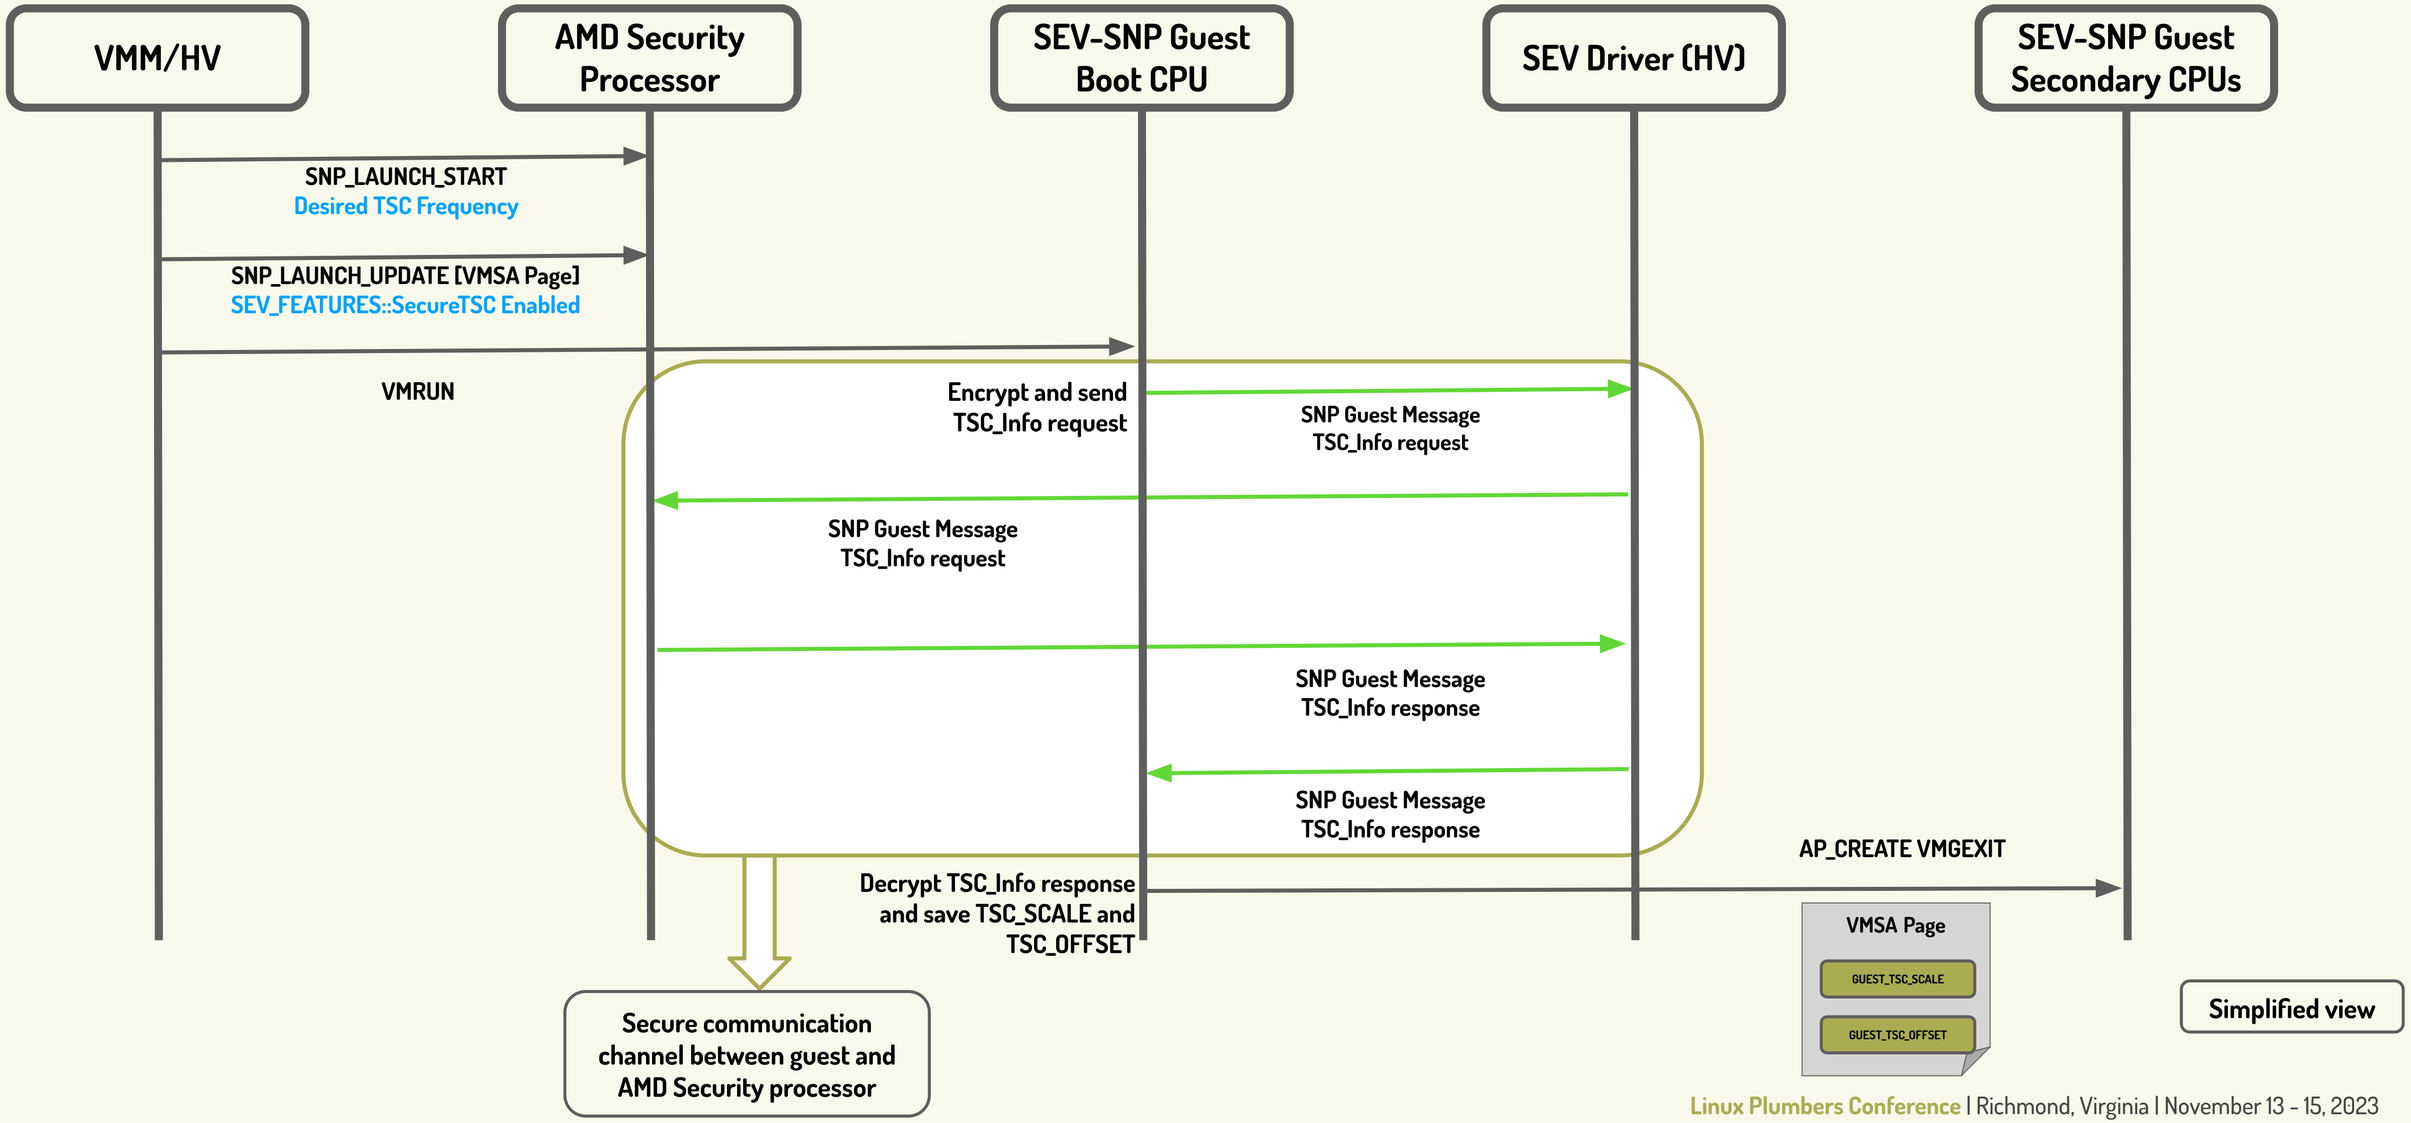

References

- SecureTSC guest v5 patches
- “Secure TSC” section in AMD64 Architecture Programmer’s Manual Volume 2: System Programming
- “TSC Info” section SEV Secure Nested Paging Firmware ABI Specification

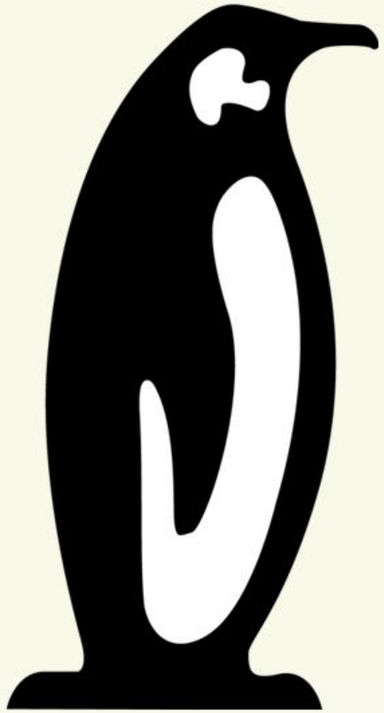

Linux Plumbers Conference

Richmond, Virginia | November 13-15, 2023

Backup

Boot sequence and sev-guest driver changes

- VMM needs to set the desired TSC Frequency during guest creation (SNP_LAUNCH_START)
- For boot CPUs GUEST_TSC_SCALE/GUEST_TSC_OFFSET is programmed by the AMD Security Processor(ASP)
- For secondary CPUs, guest needs to query TSC information from ASP Firmware
  - SNP guests have a secure communication channel to the ASP via the hypervisor
  - Guest messages are protected with an AEAD (AES-256 GCM) and sequence numbers
  - AES-256 encryption library is required for secondary CPU boot up
- SEV Guest driver provides ASP communication APIs
- Most of these functions are required during early boot for Secure TSC
- Provide clean API to SEV Guest driver and move all the core functionality in-kernel
- SecureTSC guest v5 patches: https://lore.kernel.org/lkml/20231030063652.68675-1-nikunj@amd.com/