Secure Interrupt Delivery: Lessons Learned from Alternate Injection Enablement

Dec. 12, 2025 Melody Wang

AMD
together we advance_

Secure Interrupt Delivery

- Why is secure interrupt delivery needed?
  - With SEV-SNP, the hypervisor is untrusted but in control of interrupts injected
  - Examples: recent CVEs show the hypervisor can inject malicious interrupts to break the confidentiality and integrity of the guest
    - Virtual interrupt 29 (#VC) – CVE-2024-25742
    - Virtual interrupts 0 and 14 – CVE-2024-25743
    - Int80 – CVE-2024-25744
- Solution
  - A more restricted interface between VM and hypervisor regarding interrupts
  - VM can selectively accept/drop interrupts

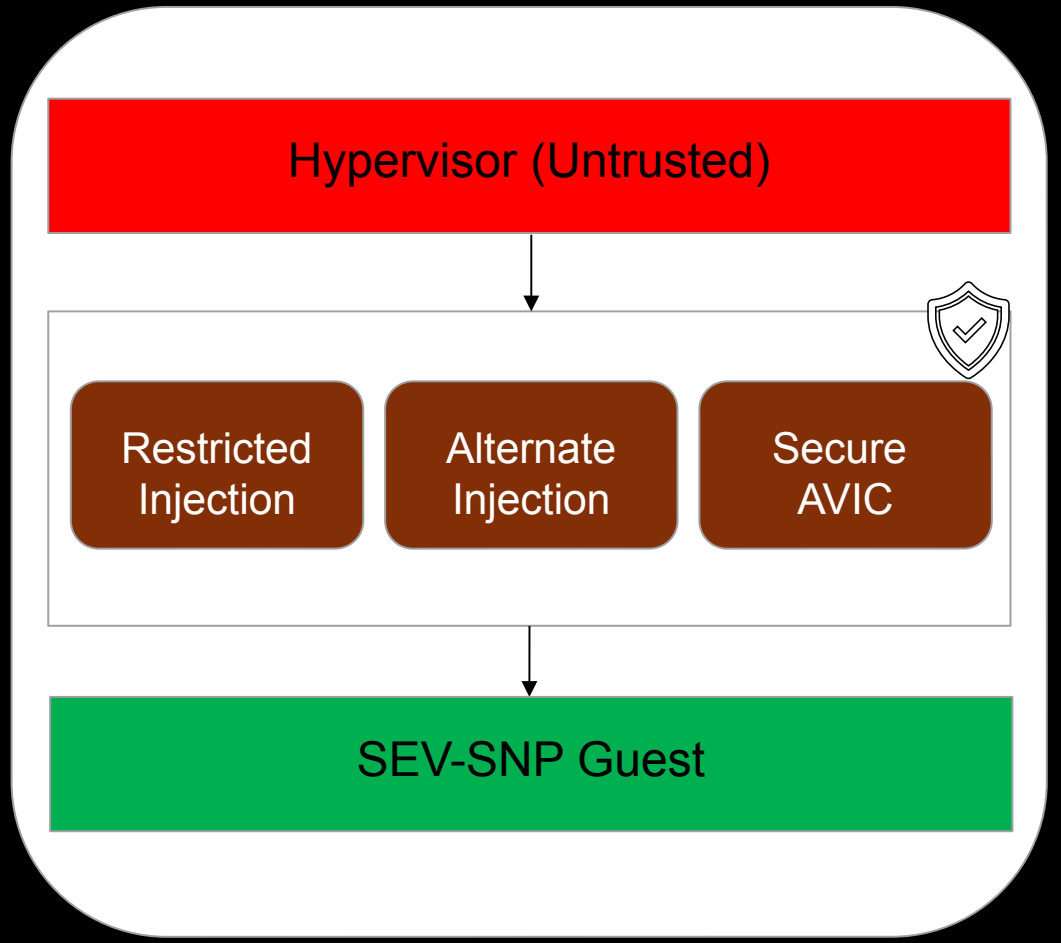

Overview of Alternate Injection

- Alternate Injection
  - Ensure interrupt presentation can only be performed by a trusted entity – the Secure VM Service Module (SVSM)
- Restricted Injection (Hypervisor -> SVSM)
  - Insulate the SVSM itself from malicious interrupts injected by the hypervisor
- Secure VM Service Module (SVSM)
  - Running at VMPL0 presents interrupts to the guest OS by writing to its VM Save Area (VMSA)

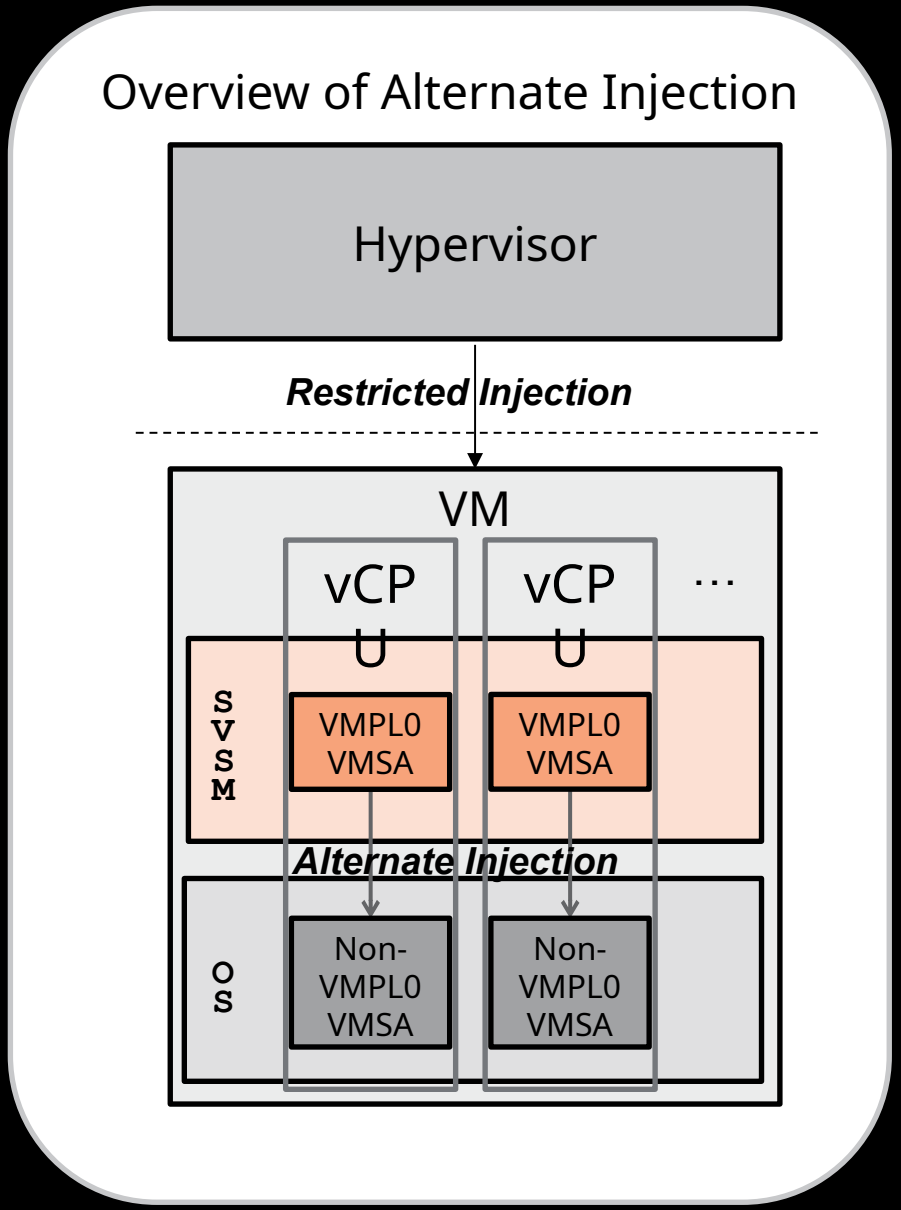

KVM Support for Alternate Injection

- Alternate Injection enabling
  - Hardware support
  - Guest request
  - Restricted Injection enabled in the SVSM
- Interrupt Injection for different VMPLs
  - Track a separate set of interrupt sources for each VMPL enabled for a given vCPU
  - Alternate Injection is not supported for VMPL0 itself
  - A new exception vector #HV
  - Target VMPL's VMSA

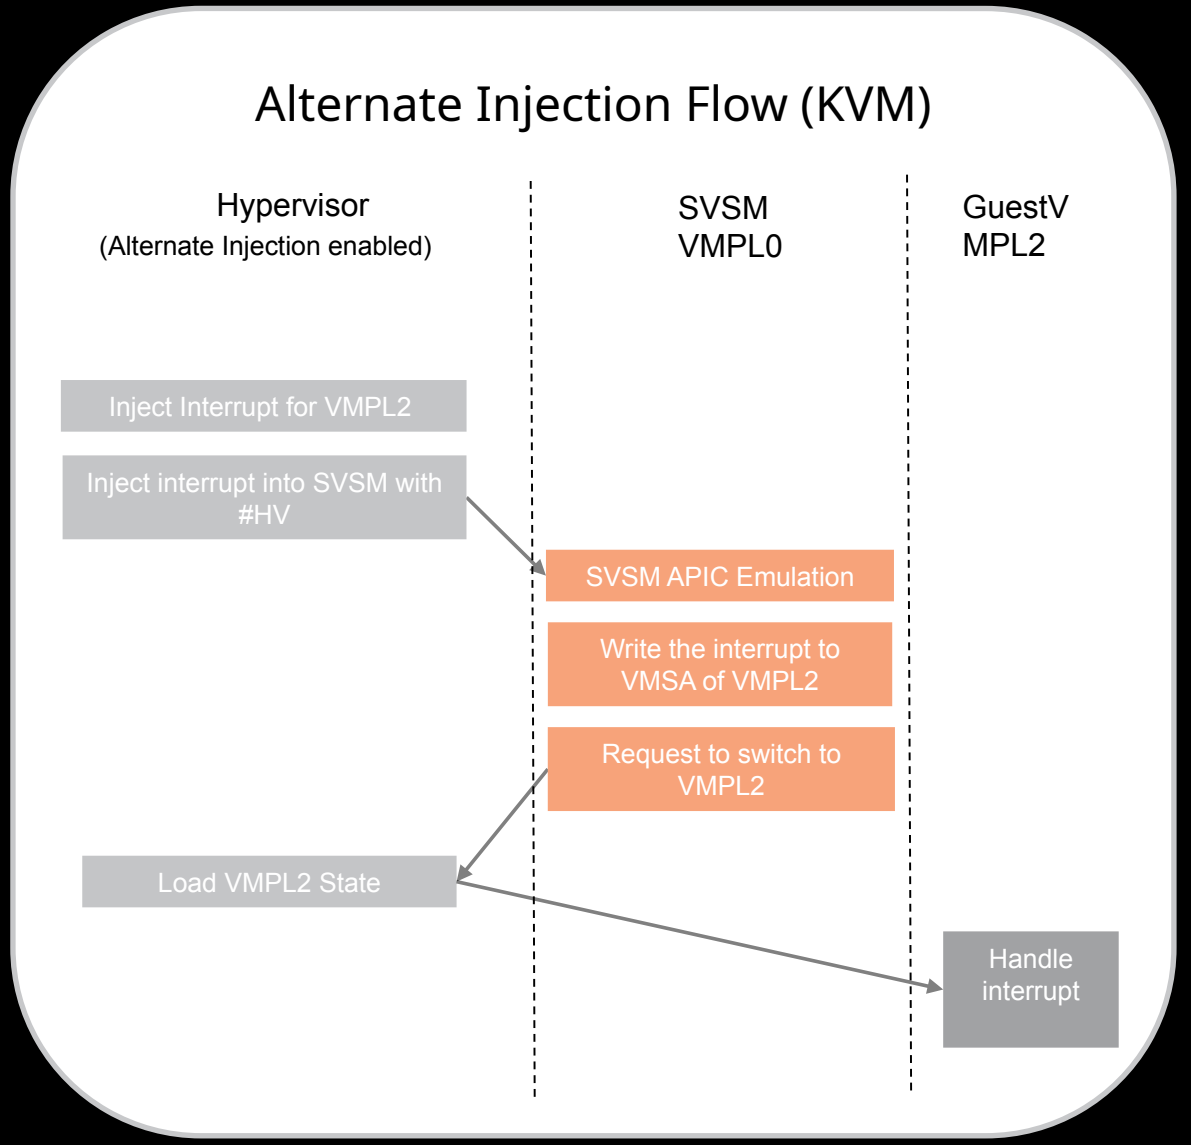

KVM design

- Design in KVM:
  - A separate path for secure interrupt delivery guests (currently implemented)
    - Avoid breaking other interrupt injection path
    - SVSM APIC used for VMPL1-3 interrupts management
    - #HV used for Restricted Injection
  - Does this seem reasonable? Other ideas?

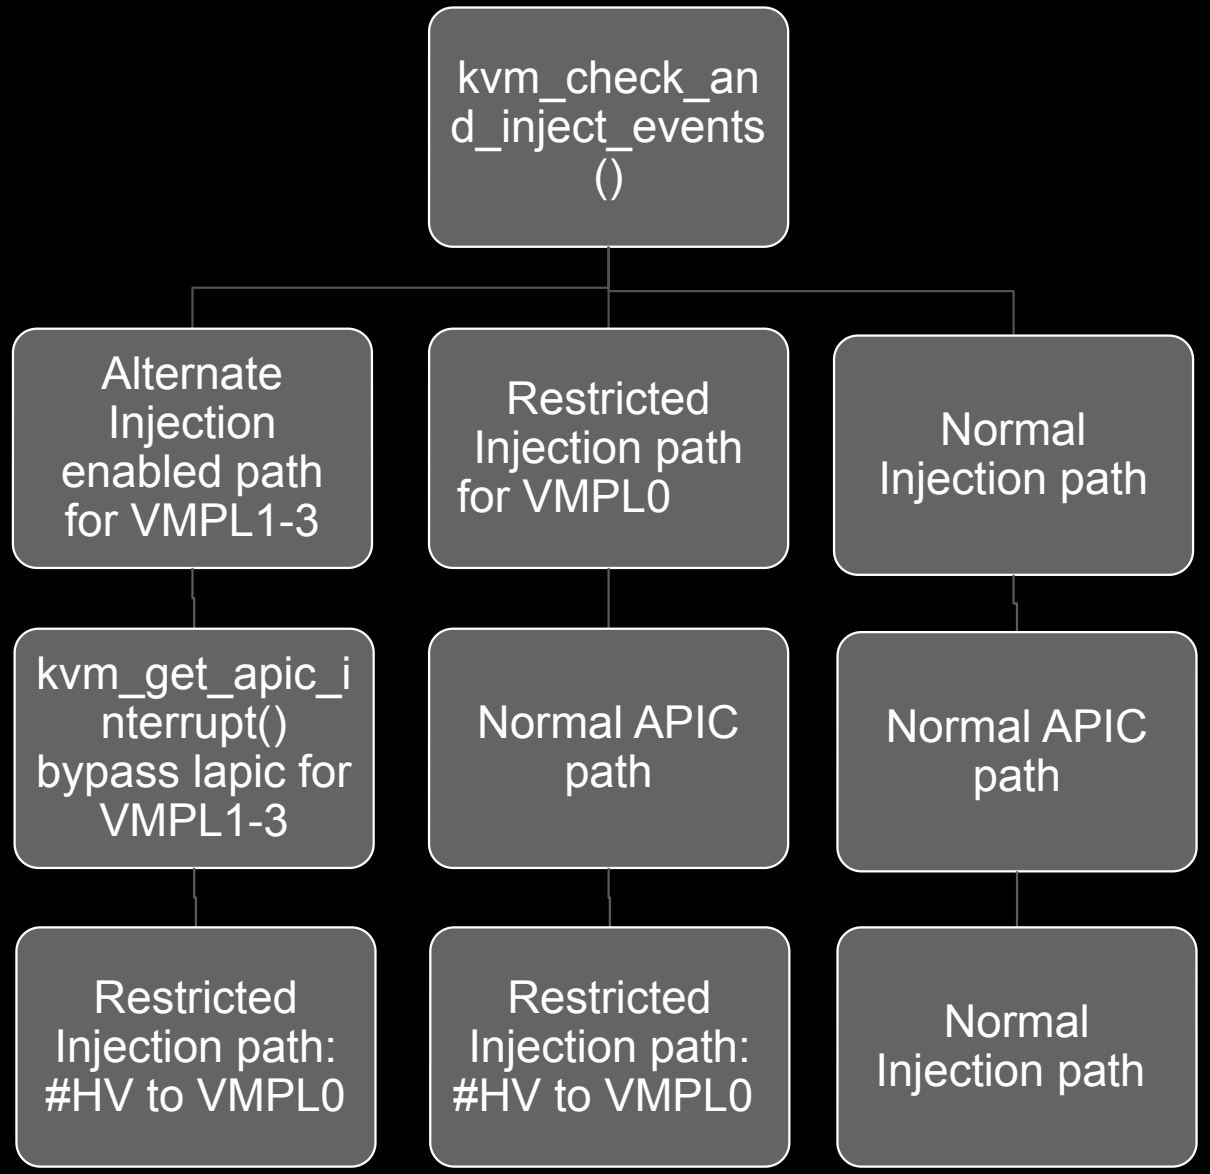

OVMF

- Problem in OVMF:
  - Starts from XAPIC while Alternate Injection only support X2APIC.
- Two possible solutions:
  - 1. Move the X2APIC enablement from the PEI to the SEC phase (currently implemented).
  - 2. Enable X2APIC mode unconditionally when running as confidential guest.

Status

- POC is done.
- Going to send OVMF, guest patches upstream.
- The current KVM patches are based on Roy Hopkins' multiple VMPL patches (https://github.com/coconut-svsm/linux, branch "svsm") and my Restricted Injection patches ([PATCH v3 0/7] Add SEV-SNP restricted injection hypervisor support - Melody Wang).
- KVM upstream patches need to be based on Planes support.


AMD

Problems (backup)

- Problem1: Interrupt disappears while being delivered to the guest.
- Root cause: An intercept happens during the interrupt delivery.

- Solution: Re-inject the interrupt when an intercept happens. Teach the SVSM to do what the hypervisor is doing.

- Problem2: Start a guest, hit an internal error – VMEXIT_BUSY.
- Root cause: Intercept happens during interrupt delivery, the busy bit gets set in the guest's VMSA.

- Solution: 1. Clear busy bit when encountering a busy bit error.
  2. Clear busy bit when EXITINTINFO.V is set.

Problems and a debug hack (backup)

- Problem1: The ICR write fails
- Root cause: Too strict restriction as to bits of trigger mode and assert in ICR message

- Solution: Ignore those bits as real hardware does

- Debugging problem: Following thousands of interrupt flows involved KVM, OVMF, the SVSM, and the guest is hard

- Solution: A debug hack integrating the SVSM, OVMF and the guest logs into KVM trace buffer by WRMSR to a non-existing MSR.