SVSM vTPM: From Boot Attestation to Persistent Storage and Beyond

Claudio Carvalho <cclaudio@ibm.com>
Stefano Garzarella <sgarzare@redhat.com>
Tyler Fanelli <tfanelli@redhat.com>

Confidential Computing MC @ LPC 2024

SVSM vTPM development status

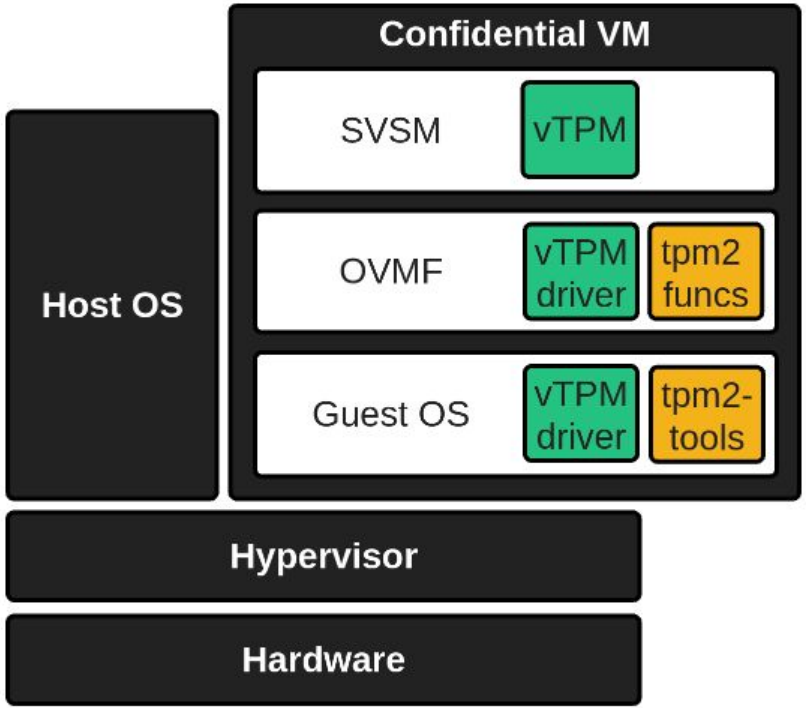

- SVSM vTPM state is **not** loaded/persisted
- SVSM vTPM is manufactured on every boot. Its seeds are randomly set.
- tpm2-tools can be used from the guest OS. E.g.:
  - Create EK: ./tpm2_createek
  - Extend PCR: ./tpm2_pcrextend

Use case: TPM-based remote attestation

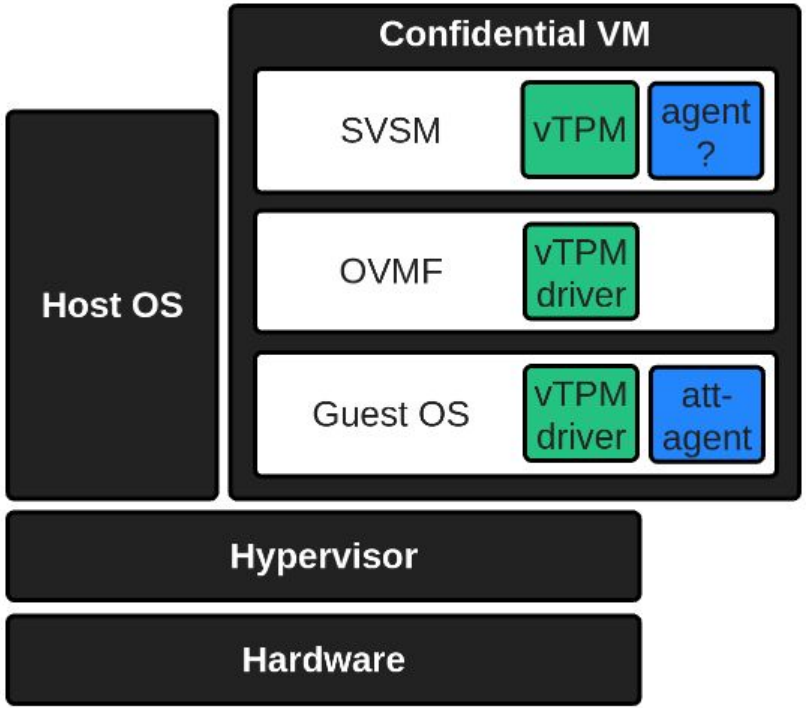

- The CVM should be attested to a known identity (vTPM EK) and state (vTPM PCRs)
- SVSM_ATTEST_SERVICE provides a VMPL0 attestation report that includes a service manifest (e.g. VTPM service). Not submitted yet.

- Injection/load of vTPM state
- Early attestation
- Persistent storage for the vTPM state
- Extend SVSM state to PCR
  - Launch measurement?
  - Event log is owned by OVMF

Use case: secure boot

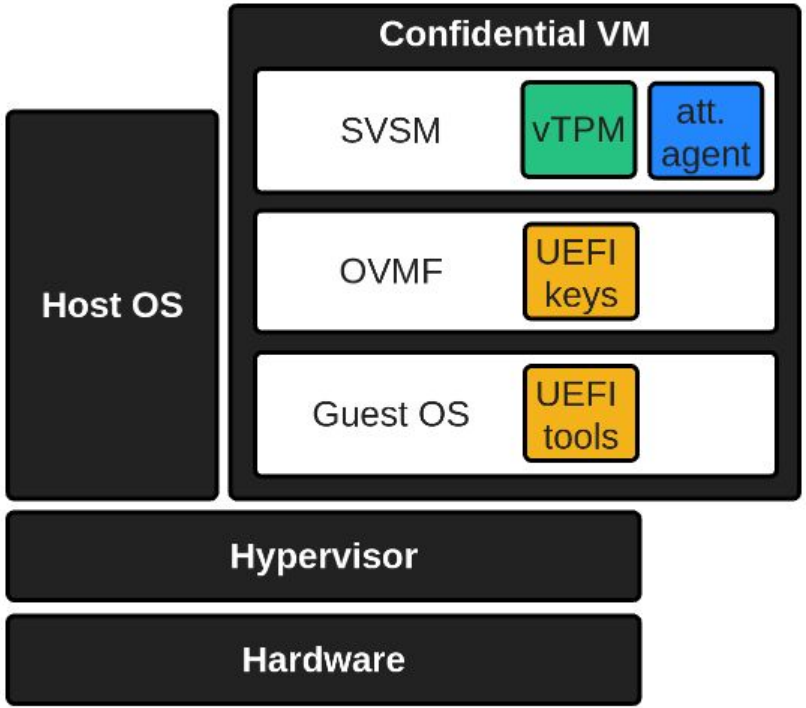

- Only firmware authorized by the signature database (UEFI DB database) can be executed
- OVMF loads the UEFI keys from a known memory region

- Inject UEFI secure boot keys
- Early attestation
- Persistent storage for the keys

Use case: TPM-based disk encryption

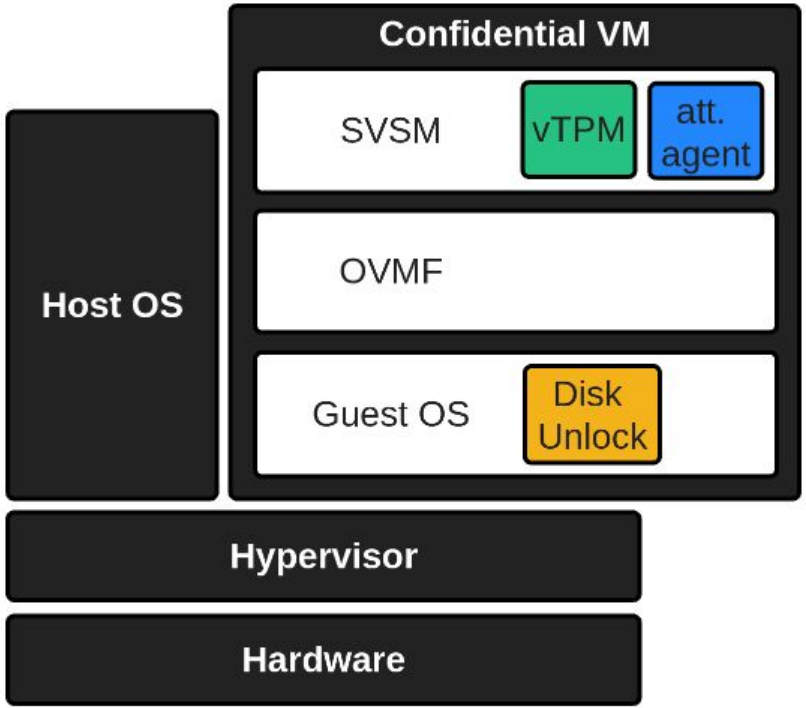

- LUKS key is decrypted using a vTPM sealed key
- vTPM key sealed to PCR

SVSM persistent state

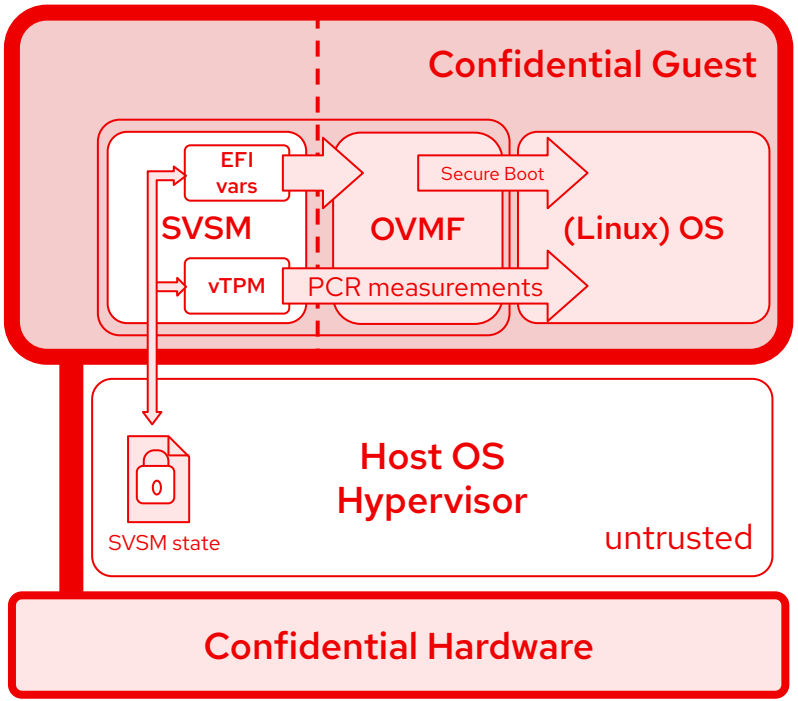

- Enable persistent vTPM and UEFI storage
  - Preserve TPM identity, counters, and storage across reboots
    - Measured boot + disk unlocking via TPM's PCR policy
  - Provide variable service in SVSM that OVMF can talk to
    - Configurable SecureBoot
- Challenges
  - Storage support in SVSM
    - Device drivers (NVRAM, virtio-blk, etc.)
    - Partitioning or simple FS
    - Integrity
    - Encryption

Rollback and clone attacks mitigation

- Attacks on the SVSM state file
  - rollback: reuse an old SVSM state
    - TPM monotonic counters could be unreliable
    - SecureBoot updates can be undone
  - clone
    - same TPM identity for different instances
- Potential mitigations
  - rollback
    - boot counter
      - released by remote attestation server
      - stored in the encrypted SVSM state
  - clone
    - only one successful attestation per boot request

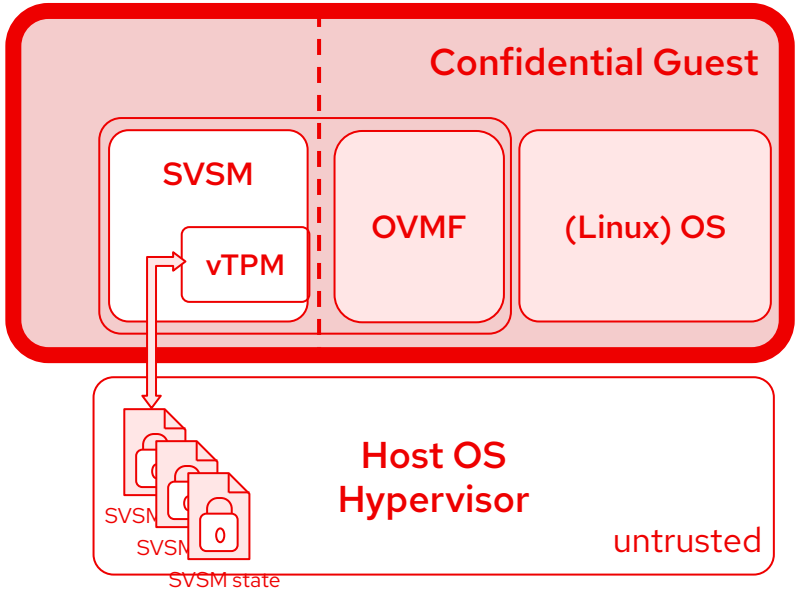

Attestation proxy

- Early attestation in SVSM is needed
  - Unlock encrypted state
  - Early Attestation and Measurement Architecture in COCONUT SVSM
- Challenges
  - No network stack
    - Proxy running in the host
  - host <-> SVSM channel
    - serial port, virtio-vsock, custom
  - Attestation protocol implementation
    - in the proxy vs in SVSM
- PoC: https://github.com/stefano-garzarella/snp-svsm-vtpm

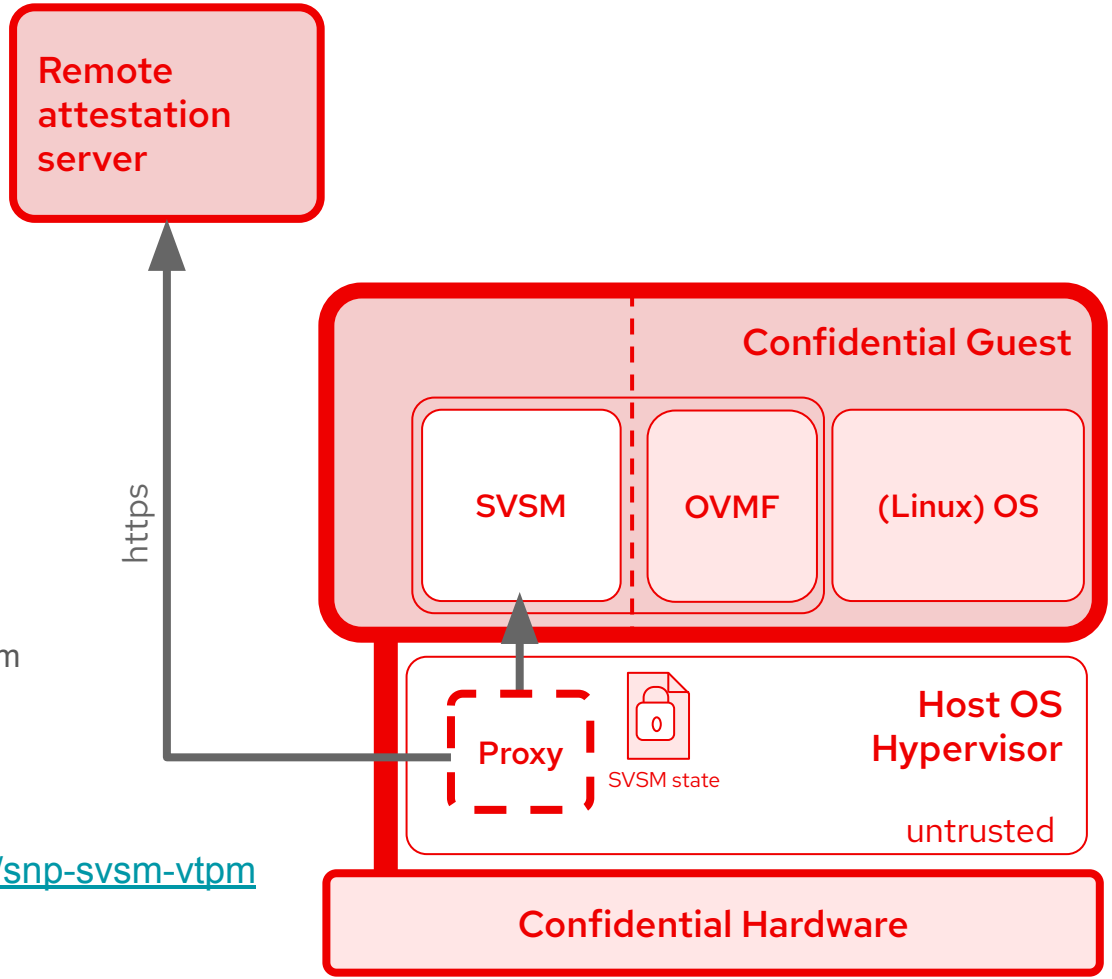

Ideas: SVSM encrypted storage

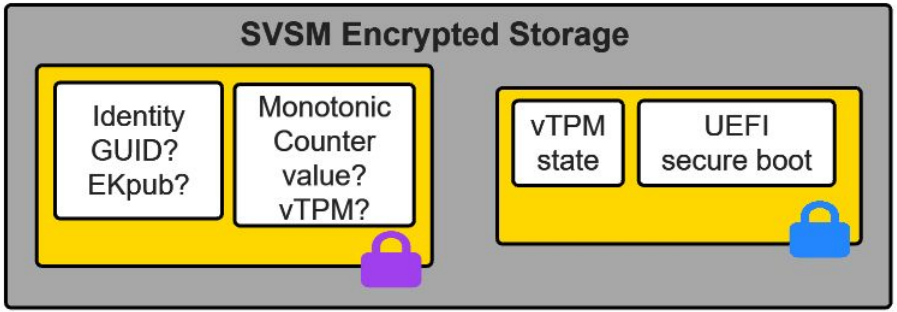

- CVM owner generates the SVSM encrypted storage
  - Manufactured vTPM
  - vTPM can contain a LUKS key sealed to a PCR
- QEMU loads the provided SVSM encrypted storage
  - Not included in the launch measurement
- Monotonic counter protects against replay attacks
- Multiple CVM instances of the same CVM could corrupt the storage

Ideas: early attestation to decrypt blue SVSM storage