Confidential VMs on public clouds and on-premise

A long way towards zero trust

Vitaly Kuznetsov

FOSDEM2025

Red Hat

What is a Confidential Virtual Machine?

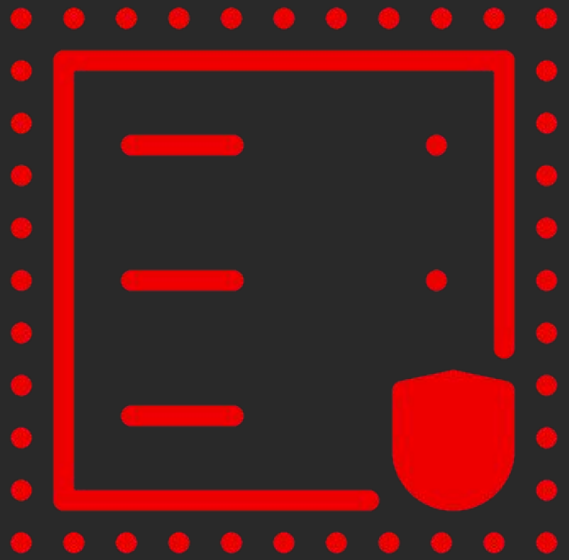

Confidentials VMs aim to provide data confidentiality: only the owner of the VM should be able to access or modify the data.

CVM cloud offerings are becoming ubiquitous

- AWS
  - SEV-SNP in GA
- Azure
  - SEV-SNP in GA
  - TDX in public preview
- Google Cloud
  - SEV-SNP in GA
  - TDX in GA

It is already possible to run CVMs on-premise

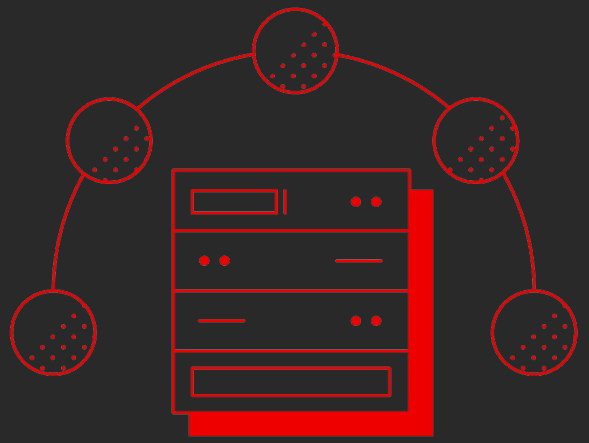

- AMD SEV-SNP
  - KVM support landed in 6.11
  - QEMU support landed in 9.1
- Intel TDX
  - Soon :-)
  - *Centos SIG for 'preview'*

So we don’t need to trust our infrastructure providers anymore, right?

Why do you trust your CVM?

(AKA "attestation")

- It is a genuine CVM running on the appropriate hardware (not an emulation!)
- The initial state of the CVM is trustworthy.
- All boot chain artifacts (bootloader[s], kernel, initramfs, ...) are trustworthy.
- The storage on the VM wasn't tampered with.
- No untrusted data injected into the VM by any provisioning agents (e.g. cloud-init).

Hardware

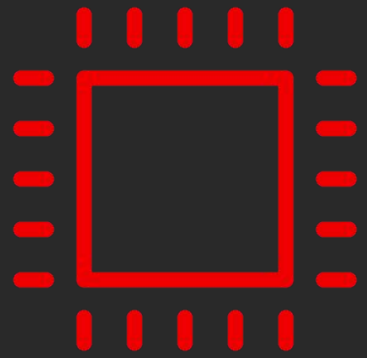

- Your root of trust is the CPU
- SNP/TDX can provide a signed report:
  - Can contain user-provided data, e.g. ssh host key to uniquely identify the VM you connect to and a timestamp to ensure freshness.
- CVMs using paravisors (e.g. Azure) may not give you raw access to the features and give you a pre-generated report stored in e.g. vTPM.
  - vTPM is your new root of trust.

Initial CVM state

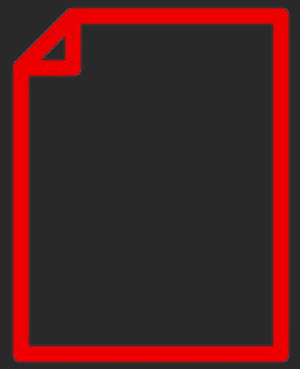

- Signed report contains “launch measurement” which describes (hash) the initial state of the memory and vCPUs.
- Firmware and instance specific data (e.g. ACPI tables) are supplied by the host
  - Can be pre-measured in the on-premise case (e.g. sev-snp-measure)
  - Can be pre-measured if you can do a reproducible build (AWS)
  - Can be pre-measured if you can bring your own firmware (Ani's talk!)
  - … or you will have to trust the opaque hash which at least doesn’t change.

Boot chain artifacts

- Launch measurements don't include binaries loaded from external storage.
- Firmware implements "vetting" (SecureBoot) or "measuring" (Measured boot) features:
  - "Vetting" is done against a varstore which may (special OVMF builds) or may not (current cloud implementations) be part of the launch measurement.
    - "Measuring" is required for the external varstore case.
- Measuring requires a TPM
  - *RTMRs can, in theory, be used on Intel TDX instead*

Measured boot: vTPMs

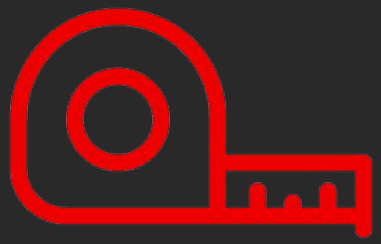

- vTPM implementation must be trusted so measurements done by it are trusted.
- SEV-SNP allows for isolated in-guest implementation using VMPLs.
  - Coconut SVSM can be used on-premise.
- Different architectures offered for Intel TDX:
  - Separate TD
  - TD partitioning
- Not always clear how cloud implementations are done.
  - Microsoft's talk in the CoCo devroom!

Storage: persistence

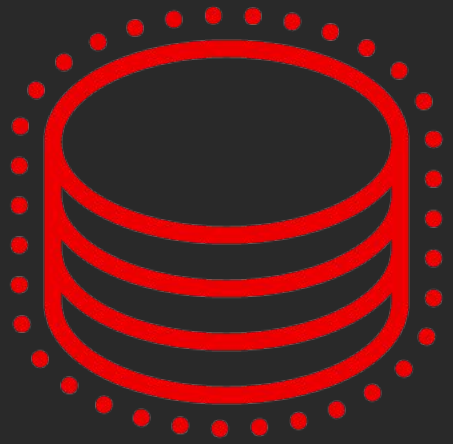

- Unified Kernel Images can help extend SecureBoot/MeasuredBoot protection to cover initramfs/cmdline.
  - My last year's talk at FOSDEM2024!
- Reliable SecureBoot db can be used for validating non-confidential, immutable parts of the storage (e.g. OS image)
- Volatile confidential storage requires encryption:
  - Key can be obtained through remote attestation
  - Key can be obtained from stateful vTPM

Stateful vTPMs

- Hyperscalers offer (AWS, GCP, Azure) stateful vTPMs for CVMs already:
  - Only Azure claims isolation of the vTPM from the host, however, the attestation process relies on other parts on Azure infrastructure.
- For on-premise deployments, Coconut SVSM is working on a stateful vTPM solution.
  - Quite complex setup, esp. for multi-tenant environments.

Provisioning agents

- No agent == no problem :-)
  - “Specialized” VM images with e.g. access keys pre-injected in the image are the best.
- Generalized VMs are at constant risk
  - All-or-nothing trust model in the existing agents.
  - No good way to authorize cloud datasource.
  - No good way to provide a custom datasource.

Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.


linkedin.com/company/red-hat


youtube.com/user/RedHatVideos


facebook.com/redhatinc


twitter.com/RedHat

Red Hat