Full disk encryption for Confidential Computing guests

OC3 2026

Emanuele Giuseppe Esposito

Vitaly Kuznetsov

Red Hat

Confidential computing

Confidential VMs ares here to stay, you can easily get one from AWS/Google/Azure/... cloud or deploy KVM + SEV-SNP/TDX on premise **today!**

- TDX, SEV-SNP reliably protect the data in use.
- Protecting the storage left to the guest operating system.

Storage protection for Confidential VMs

Must provide:

- Verity protection for readonly parts.
- Encryption + integrity protection for read-write parts.
- Attestability.
- Rollback/replay attack protection.

- **Ephemeral OS images**
  - *Confidential Containers (CoCo) is a good example.*
  - Read-only verity protected CVM root disk with optional ephemeral overlay.
  - Ship measurements that include the verity protection.
  - Encrypted data disks can be created at first boot.
- **Mutable general purpose OS images**
  - *Cloud Marketplace CVM image is a good example.*
  - Root disk must me made read-write and changes must persist.
  - Attestation server can be used only when the OS is up and running.

Verity protection for block devices in Linux

- Dm-verity is great for immutable storage!
- Rich support in systemd:
  - Tools: repart, veritysetup-generator, dissect, gpt-auto-generator.
  - Ephemeral overlay for root on dm-verity.
- Attestable:
  - The expected top level hash can be passed on the kernel command line and measured to TPM.

Creating encrypted root volumes

- Each VM instance/volume needs to be **individually** encrypted in a safe environment:
  - Pre-encrypted by a ‘trusted’ part of the infrastructure.
    - *Azure Confidential OS disk encryption is a good example.*
  - Self-encryption upon the first usage.
- Both cases require integration with attestation.

Block device encryption in Linux

- dm-crypt/LUKS is standard in Linux
- Confidential VMs add additional requirements:
  - Unique key for each volume / instance (not image!).
  - Attestable proof that the volume / master key was created in a trusted environment.
  - Confidentiality and integrity protection.
  - Rollback/replay attack protection.

Integrity protection

- Authenticated disk encryption exists but is considered **EXPERIMENTAL** in LUKSv2.
  - Provides authenticity guarantees in addition to confidentiality.
  - Additional space requirement and performance penalty.
  - Systemd (repart, dissect) support added in v260.
  - Rollback/replay protection remains a challenge.

Encryption: additional challenges

- No standard solution for storage placement randomization in Linux
  - The attacker can get **multiple** versions of ciphertext and in some cases connect it to the cleartext.
  - The attacker can restore a previous version of the sector/encryption block unnoticed.
  - The attacker can observe access patterns and thus can try to mount a side-channel attack.

Encryption: secure volume key creation

- The attacker may try to impersonate the environment, where the volume encryption key is created.
  - The proof of the encrypting environment and the measurement of the source image must be preserved.
  - See systemd upstream proposal for self-encryption.
  - Getting the encryption key from remote attestation service can also help mitigate the risk.

Combining encryption with verity

Verity -> Encryption switch for the root volume switch can provide traditional read-write OS experience.

- "Copy everything" approach: simple but inefficient.
- Use encrypted filesystem overlay: efficient for simple storage configurations.
  - Native systemd support is coming.
- Use dm-clone: possible solution for complex storage configurations.
  - Native systemd support is coming.

Almost Full Disk Encryption

- EFI system partition cannot be verity protected and/or encrypted.
- SecureBoot keys (+ Measured boot) need to be trusted for the following ESP artifacts:
  - Unified Kernel Image (UKI).
    - *Distro-shipped UKIs can be used to simplify key management.*
  - Cmdline extensions.
  - Systemd sysext/confext for initramfs extension.

TPM is always trusted

- Storage attestation story for general purpose Operating Systems fully relies on vTPM providing reliable, trusted measurements.
- Attesting vTPM itself is challenging and environment dependent.

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/@redhat


facebook.com/RedHat


x.com/RedHat

Red Hat