Full disk encryption for confidential computing guests in the cloud

Open Source Summit, Mumbai,
June 2026

Anirban Sinha
Principal Software Engineer,
Red Hat.

Red Hat

Confidential Computing

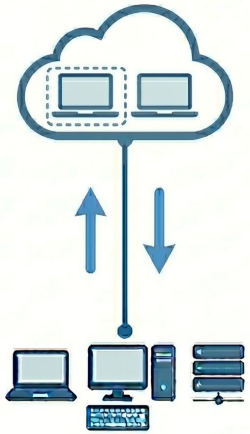

- Confidential guests are ubiquitous today both in cloud and on-prem.
- Azure/AWS/GCP and QEMU/KVM stack.
- Encrypts data in use (in memory) so that the host cannot read or modify it.
- Boot measurements guarantee deterministic and trustable boot state.

Confidential Computing

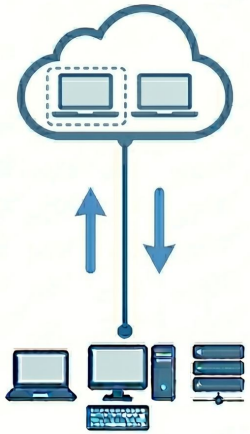

- For cloud confidential VMs, protection of data at rest (storage) is just as important.
- Disk images are managed by the cloud provider and can be easily tampered with.
- The host/provider can insert software in the disk image that can steal secrets from within the VM.

Storage protection for confidential VMs in the cloud

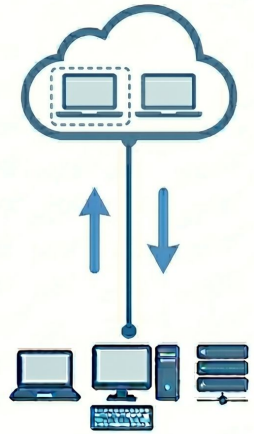

- Verity protection for readonly parts.
- Encryption + integrity protection for read-write parts.
- Attestability.
- Rollback/replay attack protection.
- We assume **vTPM is trusted.**

Cloud Disk Images

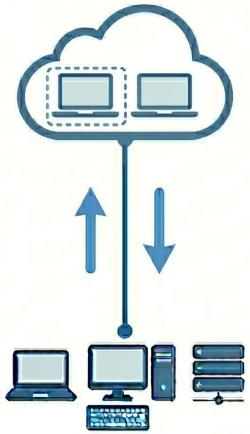

Mutable general purpose OS images

- 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.

Protecting Cloud Disk Images by Verity

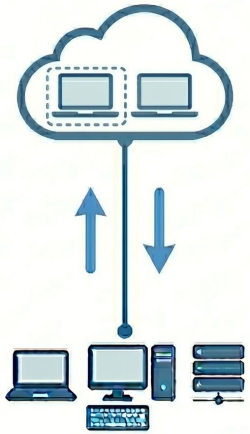

- **Dm-verity is great for immutable storage!**
- **Rich support in systemd:**
  - Tools: repart, veritysetup-generator, dissect, gpt-auto-generator.
  - Ephemeral overlay for root on dm-verity.
- **Attestable:**
  - The expected top level hash can be passed on the kernel command line and measured to TPM.
- **Not read/write by default.**

Read/Write images with encryption

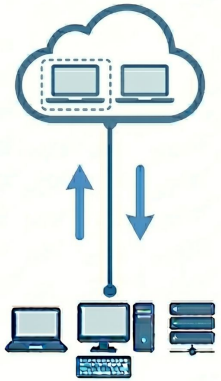

- dm-crypt/LUKS is standard in Linux

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.

Using verity protection with Encryption

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

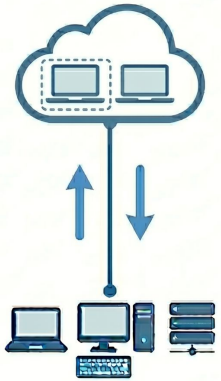

- “Copy everything” approach: simple but inefficient.
- Use encrypted filesystem overlay: efficient for simple storage configurations.
  - Native systemd support has been added via sysext/confext (in initramfs and main).
- Use dm-clone: possible solution for complex storage configurations.
  - Native systemd support is coming.

Encrypted volumes - Challenges

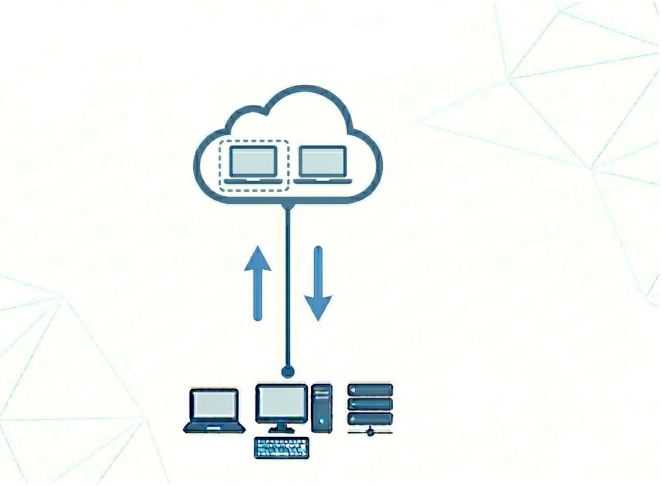

- Each VM instance/volume needs to be **individually** encrypted in a safe environment:
  - Pre-encrypted by a ‘trusted’ part of the infrastructure.
    - *Azure Confidential OS disk encryption is a good example.*
  - Self-encryption upon the first usage.
- Both cases require integration with attestation.

Encrypted volumes - 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.

Encrypted volumes - Challenges

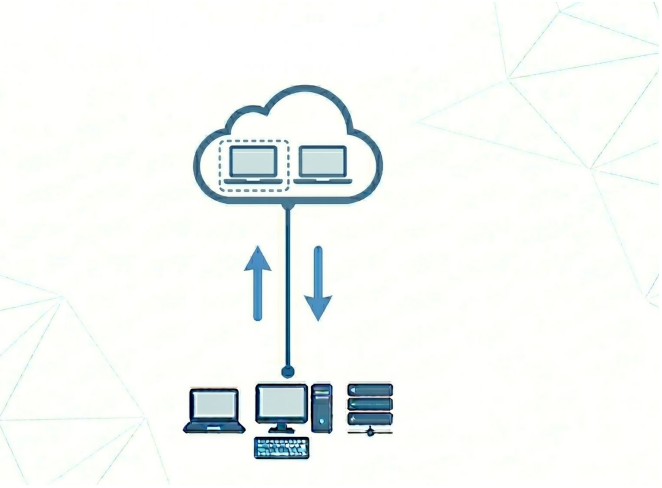

- The attacker may try to impersonate the environment, where the volume encryption key is created.
  - The proof of the encrypting environment and the measurement of the source image must be preserved.
  - See systemd upstream proposal for self-encryption.
  - Getting the encryption key from remote attestation service can also help mitigate the risk.
    - ⇒ Systemd-repart with EncryptToken option is proposed here for LUKS v2.

Encrypted volumes - Integrity protection

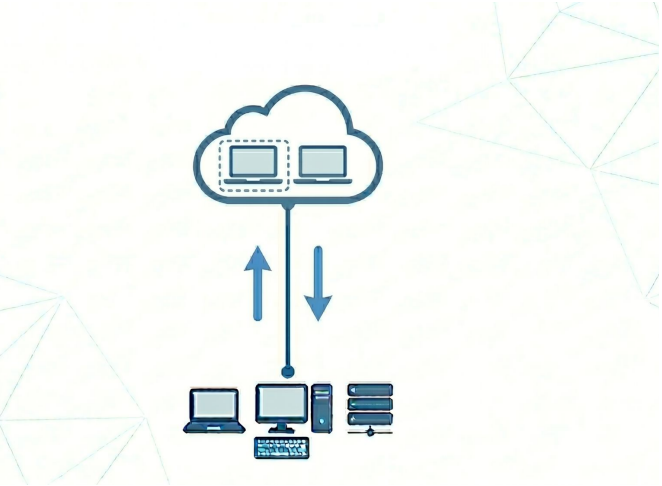

- Authenticated disk encryption exists but is considered **EXPERIMENTAL** in LUKSv2.
  - Provides authenticity guarantees in addition to confidentiality.
  - Additional space requirement and performance penalty.
  - Systemd (repart, dissect) support added in v260.
  - Rollback/replay protection remains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

The boot EFI partition

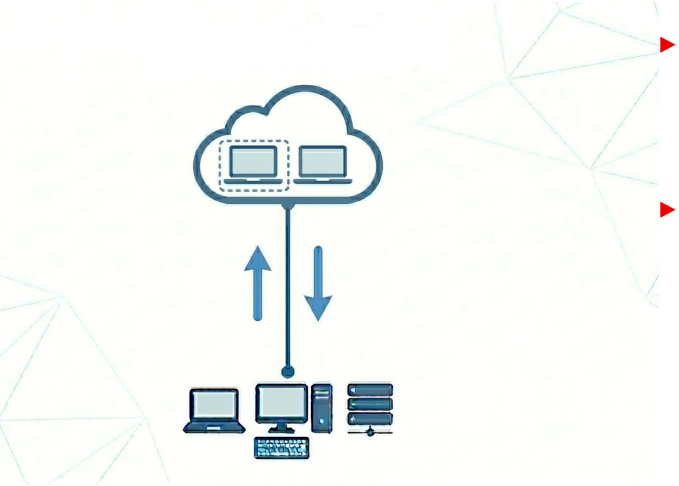

- EFI system partition cannot be verity protected and/or encrypted.
- SecureBoot keys (+ Measured boot) with signed UKIs can be used to validate the boot chain.

Closing thoughts

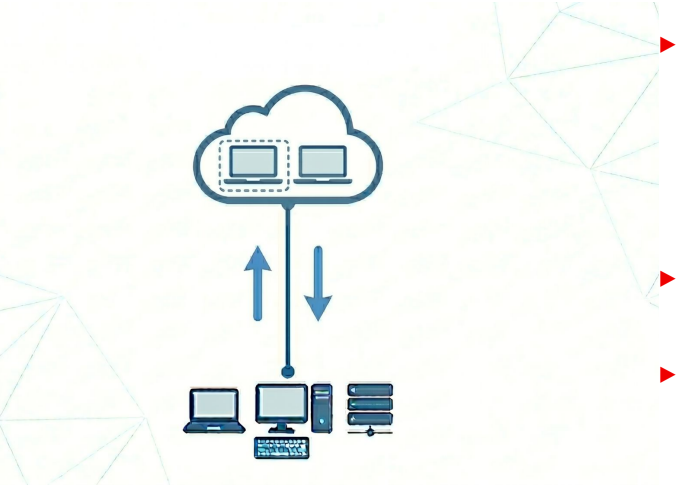

- Confidentiality and integrity of storage data for confidential cloud deployments is just as important as confidentiality of data in memory.
- There is no off-the-table solution.
- Multiple open source Linux solutions exists and some need to be added that can be combined to get the desired confidentiality of data at rest.
- The problem is non-trivial to solve.

Things to read and do

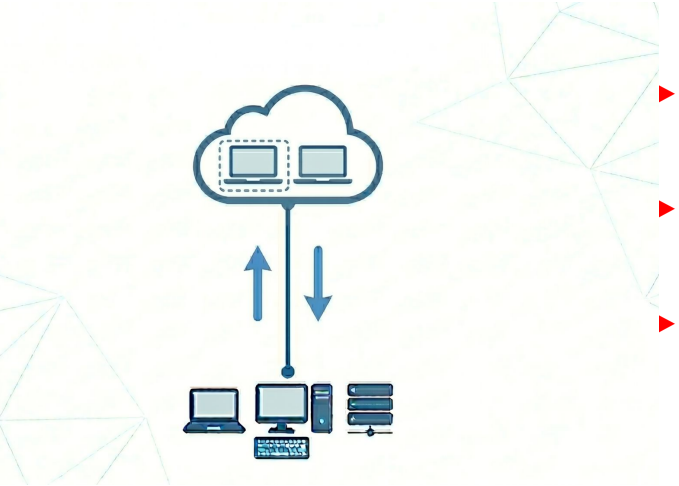

- Confidential computing primer
- Confidential virtual machine storage attack scenarios
- How to encrypt RHEL images for Azure confidential VMs
- Cool demos around vTPM and confidential VMs.

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

facebook.com/redhat

youtube.com/@redhat

x.com/RedHat

Red Hat