Running Confidential Workloads with Podman

Sergio Lopez
slp@redhat.com

A (very) brief introduction to Confidential Computing

As defined by the Confidential Computing Consortium

"Confidential Computing is the protection of data in use by performing computation in a hardware-based, attested Trusted Execution Environment."

Where we are today...

Virtualization-based Confidential Computing is available since 2017

Support for SEV and SEV-ES is merged in mainline Linux and shipped by most distributions

But barely anyone is using it :-(

But why?

Doing Confidential Computing “The Right Way” it’s complicated.

What the hardware gives us is just the primitives, without a defined way how to use them.

What do you need to measure, how and when are questions that can heavily depend on the context.

The Idea

Integrate Confidential Computing in the Container workflow

podman + crun + libkrun

The Idea

Integrate Confidential Computing in the Container workflow

podman + crun + libkrun

A Virtual Machine Monitor written in Rust, provided as a dynamic library, bundled with a minimalist Linux kernel (and FW and initrd), that's able to use Confidential Computing tech

What's a Confidential Workload? (I)

The Goals

- **Must be compatible with the existing container tools and workflows**
  - Deployed as an OCI
  - All the information needed to run as a TEE must be self-contained

- **Must meet the Confidential Computing requirements**
  - The disk must be encrypted and integrity protected
  - The initial memory contents must be easy to measure
  - Host leaks must be limited, even if that means breaking some of the conventional container semantics

What's a Confidential Workload? (II)

The result

- **A regular OCI image with the following contents**
  - TEE-specific parameters
  - A LUKS-encrypted disk image with the contents of another OCI

- **A Virtualization-based VM/TEE provided by libkrun**
  - Well-known set of initial memory contents (provided by libkrunfw)
  - No host leaks allowed, except through the network
  - Memory encryption, integrity protection and attestation by using Confidential Computing HW (SEV, SEV-ES, SEV-SNP, and soon TDX).

The execution context

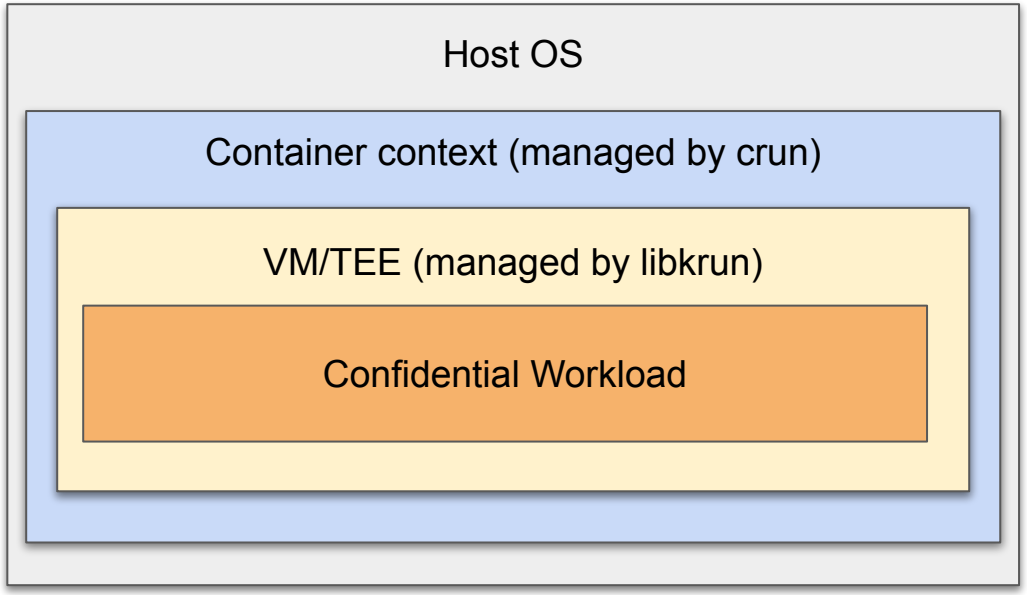

Red Hat

How is each data context protected

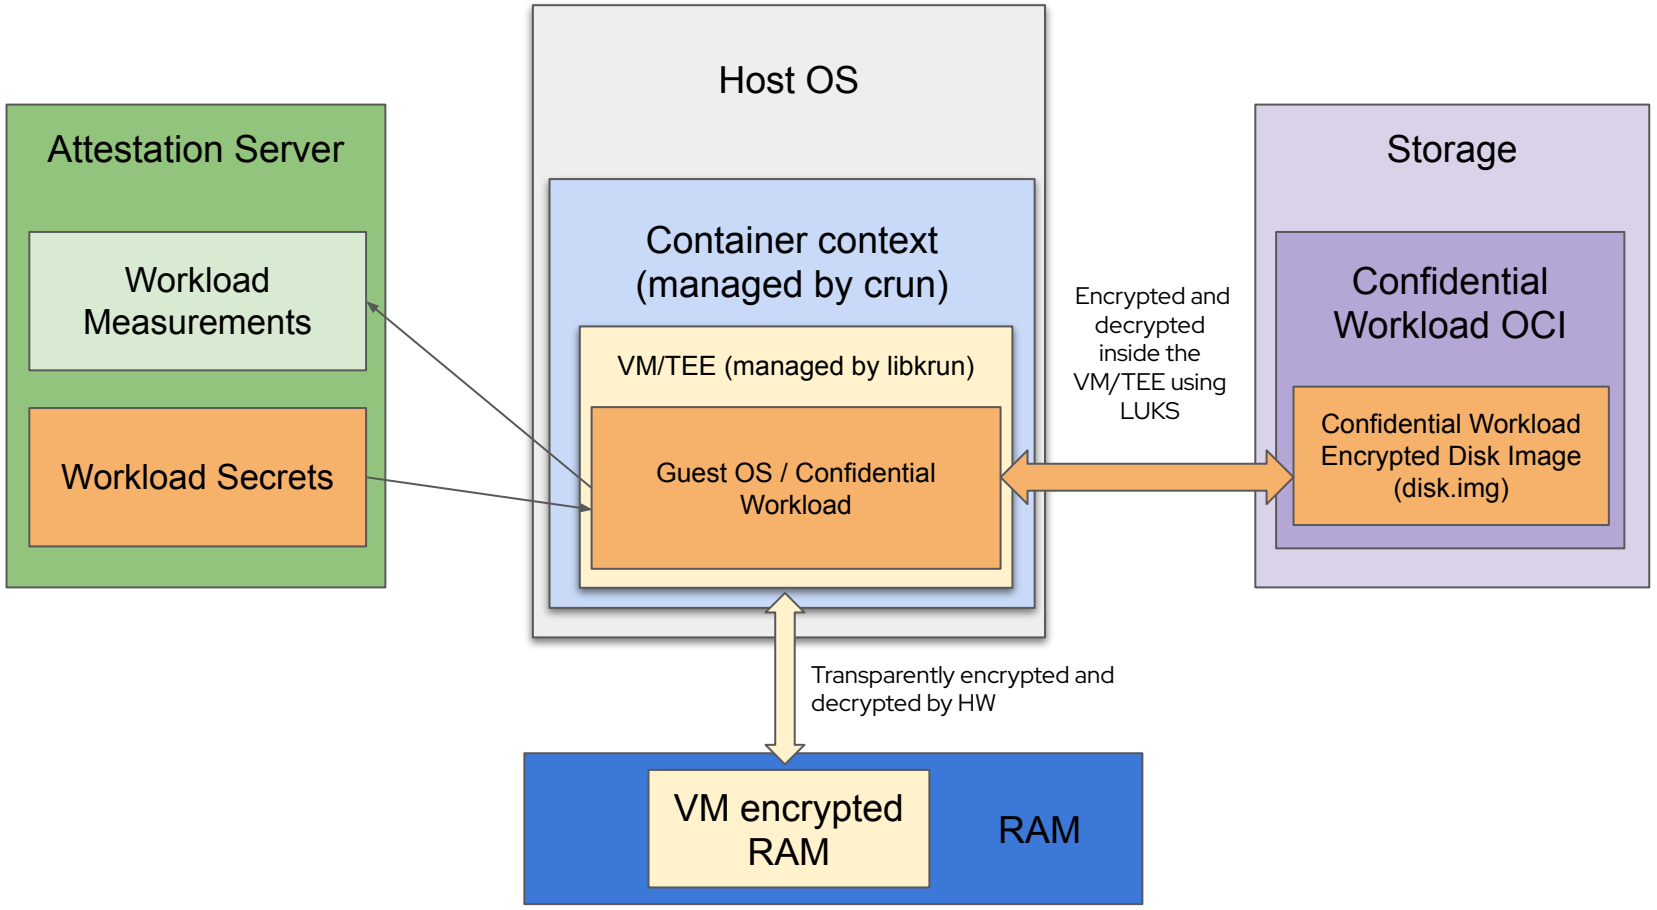

Red Hat

From OCI to Confidential Workload

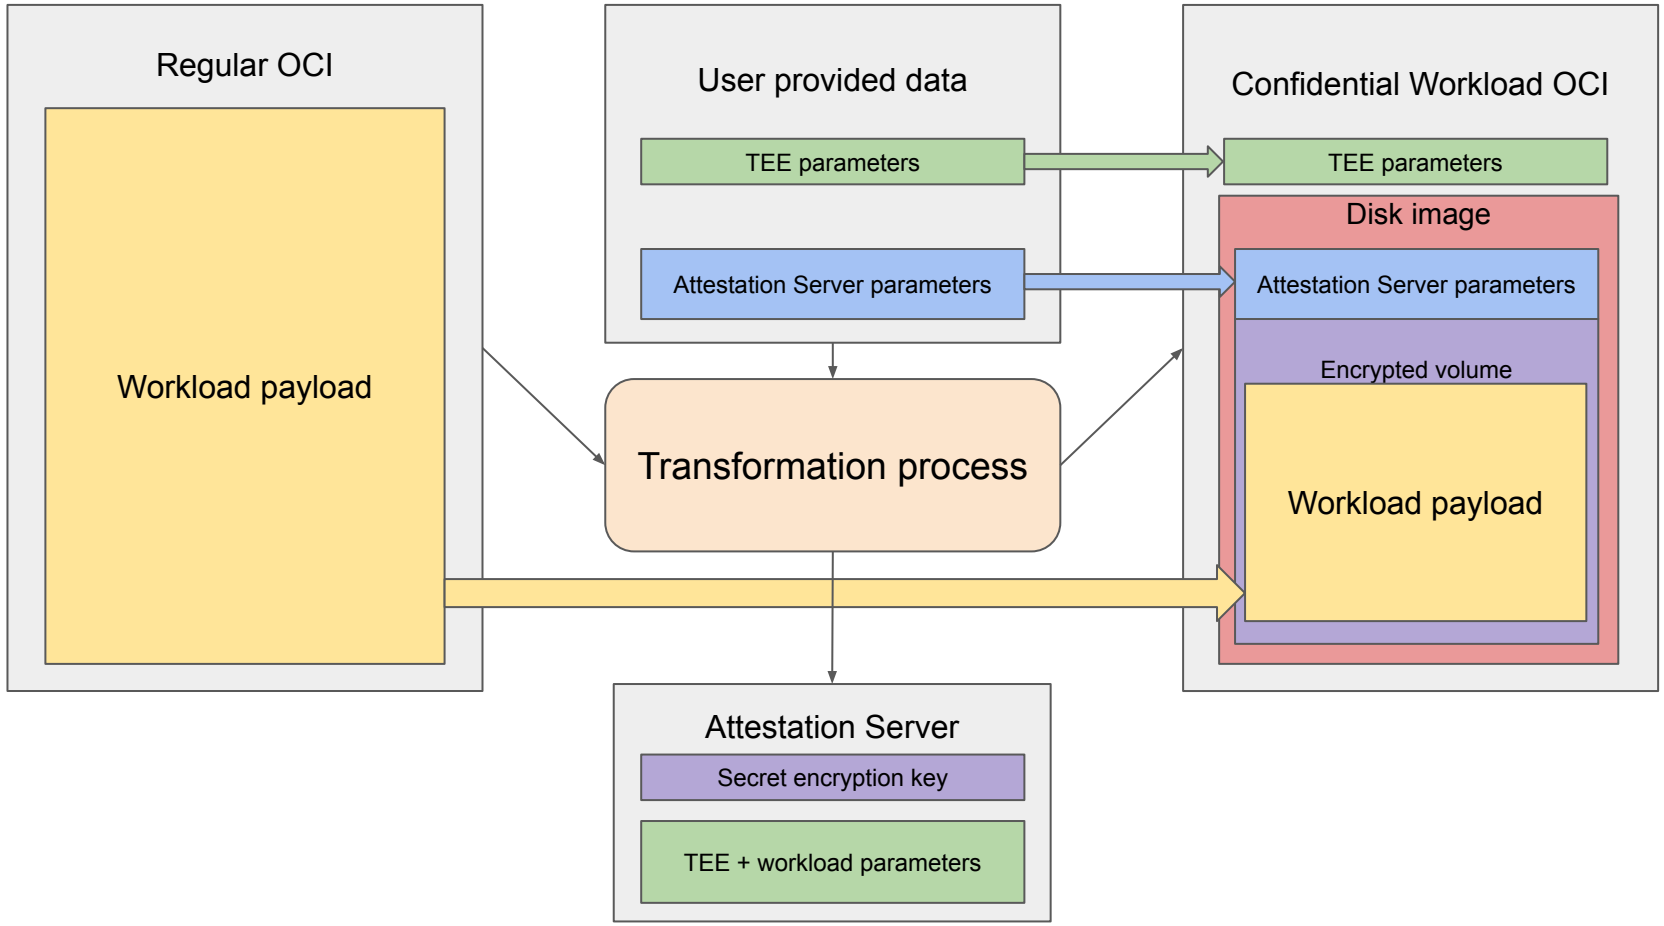

Red Hat

Confidential Workloads vs. Kata Confidential Containers

The match of the century!

podman + crun + libkrun

vs.

CONFIDENTIAL CONTAINERS

Confidential Workloads **&** Kata Confidential Containers

Why not both?

- **When to use Kata Confidential Containers**
  - Migrating existing container deployments
  - Cloud deployments with many containers per pod

- **When to use Confidential Workloads**
  - Cloud deployments with many single-container pods (i.e. FaaS)
  - Non-cloud deployments (Embedded, Edge, Automotive…)

DEMO