Contribution ID: **25** Type: **not specified**

Running Confidential Containers

Tuesday, 21 September 2021 08:55 (25 minutes)

Nowadays, containers are a private and public cloud commodity. Isolating and protecting containerized workloads not only from each other but also from the infrastructure owner is becoming a necessity.

In this presentation we will describe how we're planning to use confidential computing hardware implementations to build a confidential containers software stack. By combining the hardware encryption and attestation capabilities that these new ISAs provide, the proposed software architecture aims at protecting both container data (downloaded from container image registries and generated at runtime) and code from being seen or modified by cloud providers and owners.

As the Kata Containers project already uses hardware virtualization to provide a stronger container isolation layer, we will first explain why and how we want to use the Kata runtime and agent as the foundation for running confidential containers.

Then we will look at the container specific requirements that we need to take into account for building that software stack. Short boot times, low memory footprint or the inherently dynamic, ephemeral and asynchronous nature of container workloads are some of the technical challenges that we're facing when it comes to running confidential containers. The final part of the talk will go through some of the technical solutions we're building to address those challenges. In particular we will speak about:

- Transparent memory and cpu state encryption: As the Kata runtime can run on top of heterogeneous nodes, running different confidential computing implementations (TDX, SEV, etc), we have to build a small framework for transparently enabling the underlying encryption technologie whenever a confidential container is scheduled on a given node.
- Container image service offload: The entire container software ecosystem assumes container image layers can be downloaded, unpacked and mounted from the host itself. This obviously breaks the confidential computing security model and that brings the need for offloading at least part of the container image management from the host to the guest.
- User space attestation: Adding container images to the initial guest image and measurements can have a large impact on boot time and also break the dynamic and ephemeral nature of container workloads in e.g. a Kubernetes container orchestration context. As a consequence container images must be downloaded, and either verified or decrypted from the guest user space Kata agent, who then becomes responsible for triggering the container image credential provisioning through e.g. remote attestation.

I agree to abide by the anti-harassment policy

I agree

Primary author: ORTIZ, Samuel

Presenter: ORTIZ, Samuel

Session Classification: Confidential Computing MC

Track Classification: Confidential Computing MC