Contribution ID: **239** Type: **not specified**

Guest private memory for software-based hypervisors

Thursday, 19 September 2024 11:10 (20 minutes)

This talk presents different proposals for supporting guest private memory in Android for Arm64 in the pKVM and the Gunyah hypervisors.

Confidential computing is gaining popularity, with hardware-based (Intel TDX, AMD SEV, Arm CCA) and software-based (pKVM, Gunyah) solutions. A common aspect is the ability to create a "protected" guest, whose data is neither inaccessible by other VMs nor by the host itself, unless explicitly shared by the guest.

In the original KVM API, guest memory is provided as a host user space address to KVM, and is mapped by the host. Although the hypervisor prevents the host from accessing the guest memory via that address, an erroneous access could be fatal to the system and result in a full reset.

To address these issues, guest_memfd() was created as a new API. It represents guest memory using a file descriptor, along with an allocator that restricts what can be done with that memory, such as mapping it at the host. With the guest memory not being mappable to begin with, erroneous accesses cannot take place.

The pKVM and the Gunyah hypervisors target mainly Android on Arm64. They use hypervisor (stage 2) page table protection, not encryption, to protect guest memory. Among other things, this allows in-place guest memory conversion between shared and private. However, the current guest_memfd() implementation never allows guest memory mapping, and sharing is done by copying the data

In this talk, we propose modifications to guest_memfd(), as well as alternative approaches, to enable these hypervisors to perform shared to private conversions in-place (and vice versa).

So far, we have presented two proposals as RFCs upstream [1, 2], followed by discussions on the best approach moving forward. This talk aims to summarize these discussions to reach a solution consistent with existing approaches.

[1] https://lore.kernel.org/all/20240222161047.402609-1-tabba@google.com/
[2] https://lore.kernel.org/all/20240618-exclusive-gup-v1-0-30472a19c5d1@quicinc.com/

Primary authors: BERMAN, Elliot (Qualcomm); TABBA, Fuad (Google)

Presenter: TABBA, Fuad (Google)

Session Classification: KVM Microconference

Track Classification: KVM Microconference