Hadivoli

Fine-grained process isolation against hostile OS and libraries using Confidential VMs

Harena Rakotondratsima[1], Jean-François Dumollard[1], Jana Toljaga[1], Nicolas Derumigny[2], Gaël Thomas[2]
1 . firstname.lastname@telecom-sudparis.eu 2 . firstname.lastname@inria.fr

Virtual Machine (VM)

- **Hardware abstraction layer** technique
- Multiple **isolated systems** on one physical machine
- **Hypervisor:** controls and manages VMs
- **VMs**: guest system executing with dedicated **virtual resources**

Confidential Virtual Machines (CVMs)

- VMs run inside **encrypted memory**
- Protects VMs against **external threats:**
  - hypervisor
  - physical attacks (bus snooping)
  - malicious admins

Applications in VMs are exposed to **internal threats**:

⚠ Kernel requires **full memory access**

⚠ **Untrusted** third-party libraries

=> **Large attack surface**

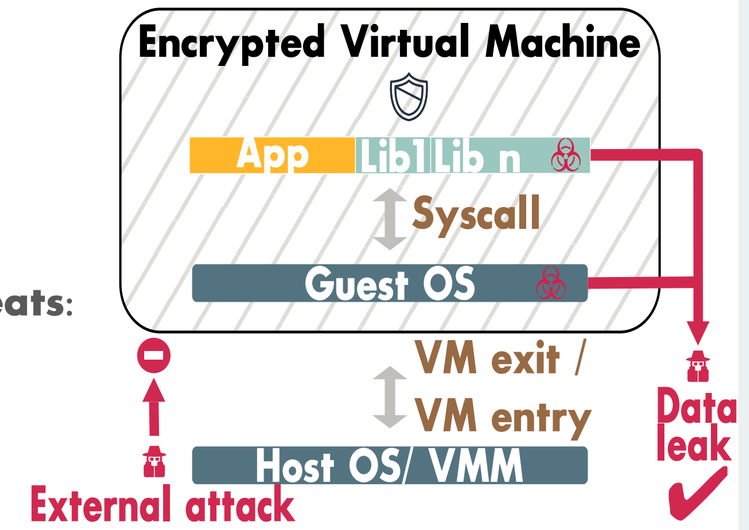

Challenge

How to prevent internal threats from leaking application data?

Libvolimem: light virtual machine

- **Process virtualization** without a full OS
- **Runtime** with 8K LOC
- **Kernel capabilities** exposed to user space
- Overhead induced by transfers (e.g., syscalls, shared memory) and virtualization

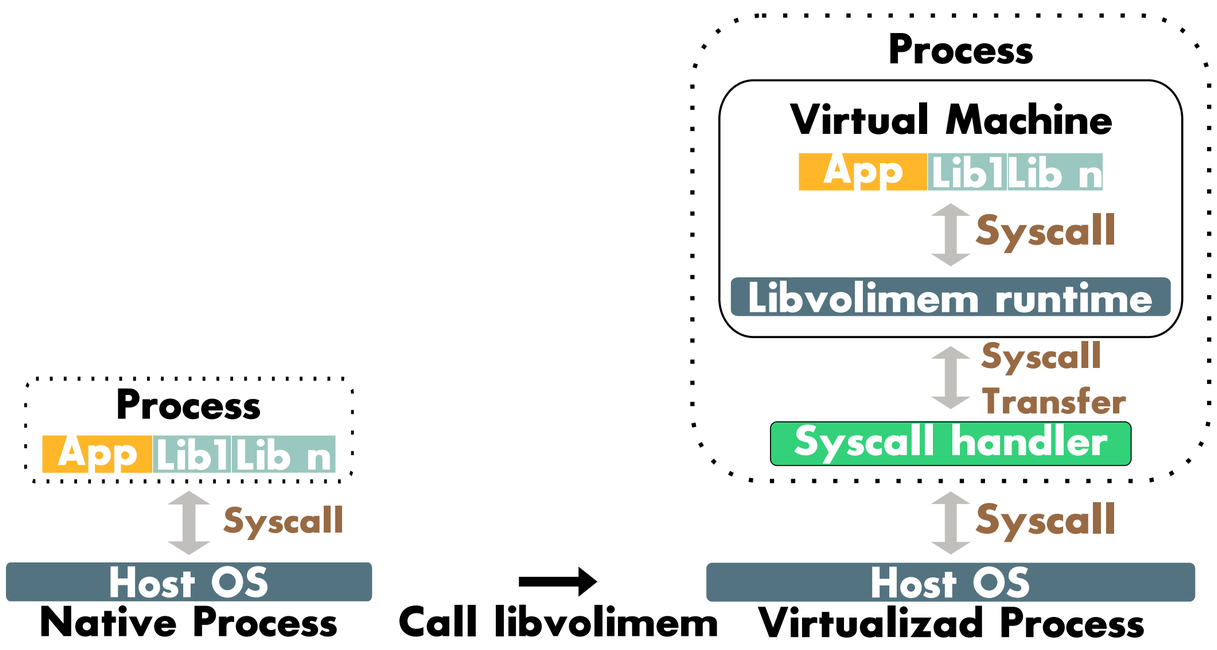

Out of scope

- Denial-of-service attacks by the host OS or untrusted libraries
- Logical vulnerabilities or bugs inside the trusted application
- Side-channel attacks (cache-timing attacks, ...)

Expected results

✔ **Reduced attack surface:** fewer code paths and potential exploitation vectors

✔ **Process level isolation:** protection from external and internal threats

✔ **Flexible trust boundary:** the application owner decides which libraries to trust

Hadivoli: process virtualization with Intel TDX

- Run the process inside an **Intel TDX-encrypted VM:**
- **Trusted** context: application + libvolimem
- **Untrusted** context: third-party libraries
- **Memory isolation** between contexts
- **Secure channel** communication between contexts
- **Syscall transfer** using **unencrypted shared** buffer
- **Mediates** untrusted context syscalls to prevent data leaks

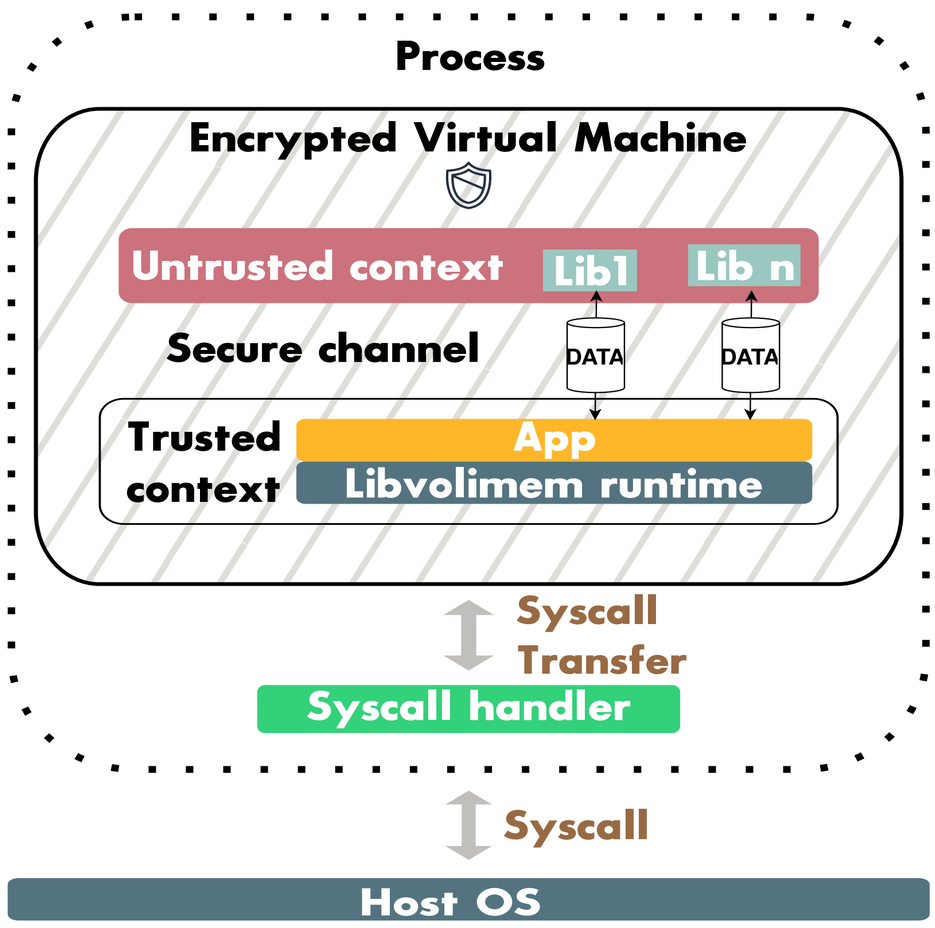

contact: harena.rakotondratsima@telecom-sudparis.eu