VEIL: A Protected Services Framework for Confidential Virtual Machines

Adil Ahmad[1], Botong Ou[2], Congyu Liu[2], Xiaokuan Zhang[3], Pedro Fonseca[2]

1 2 3

CVMs restrict data access from outside components

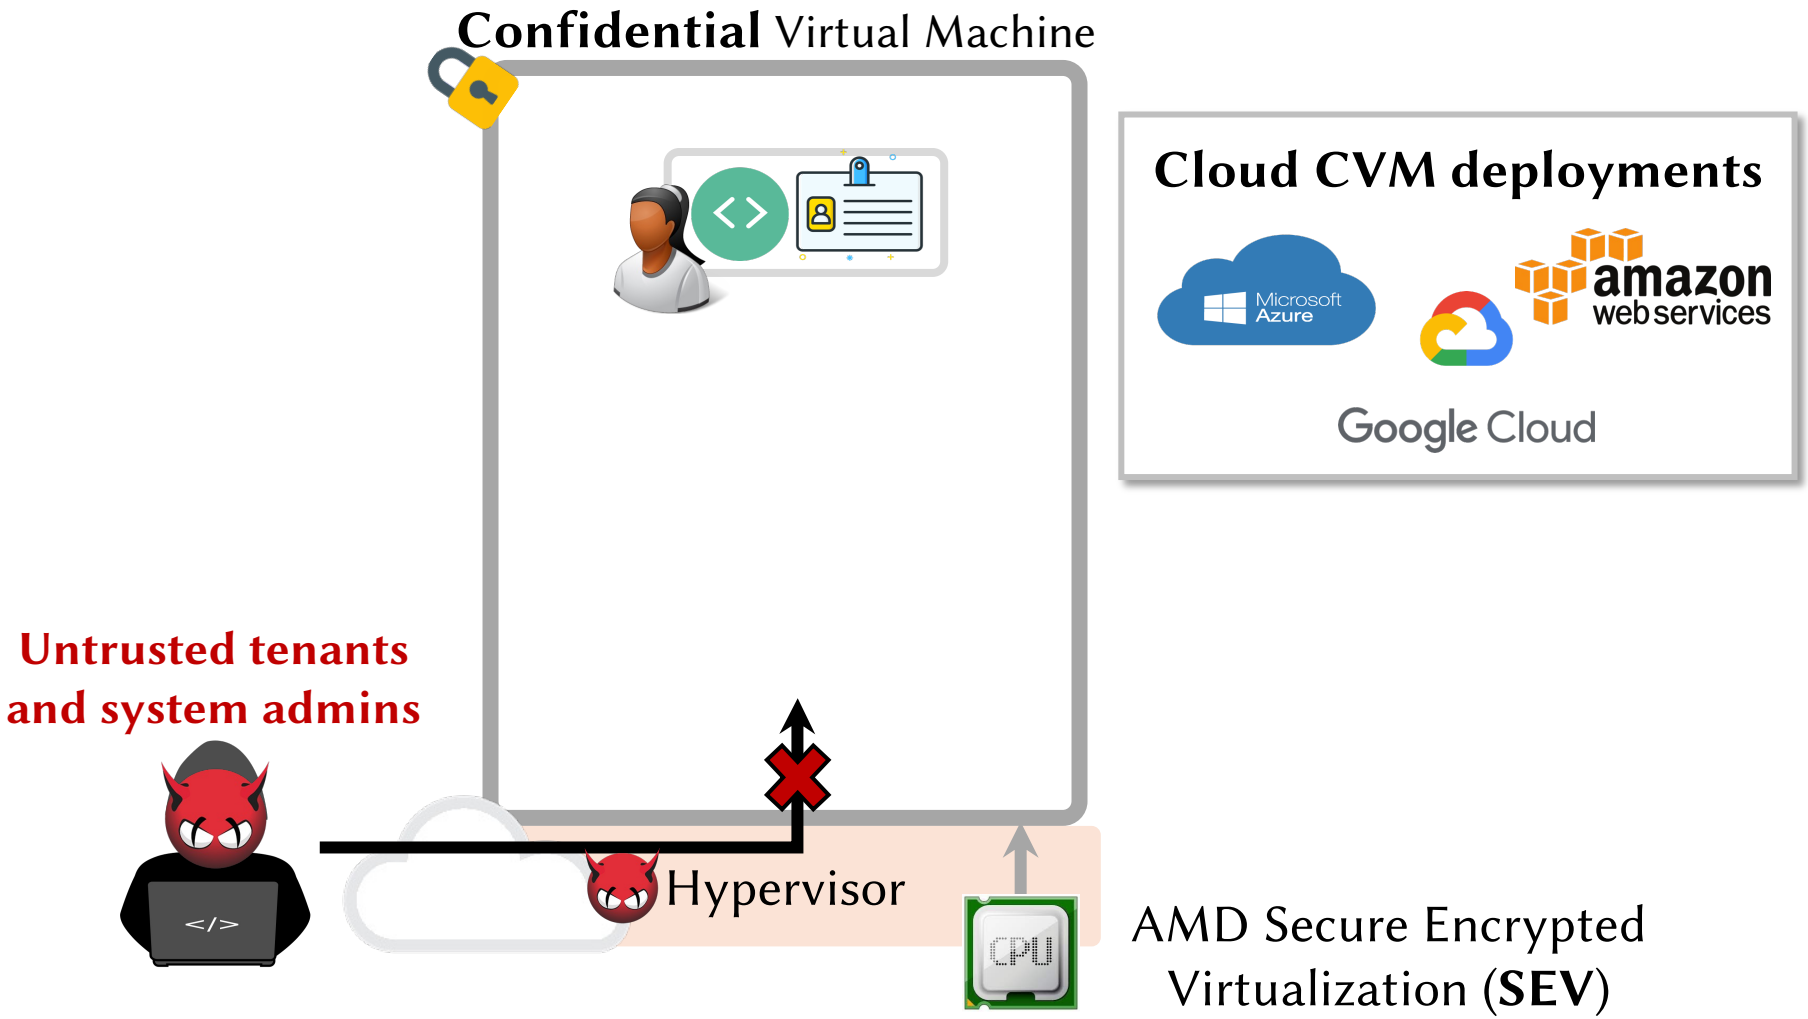

CVMs rely on OS for defense against remaining attacks

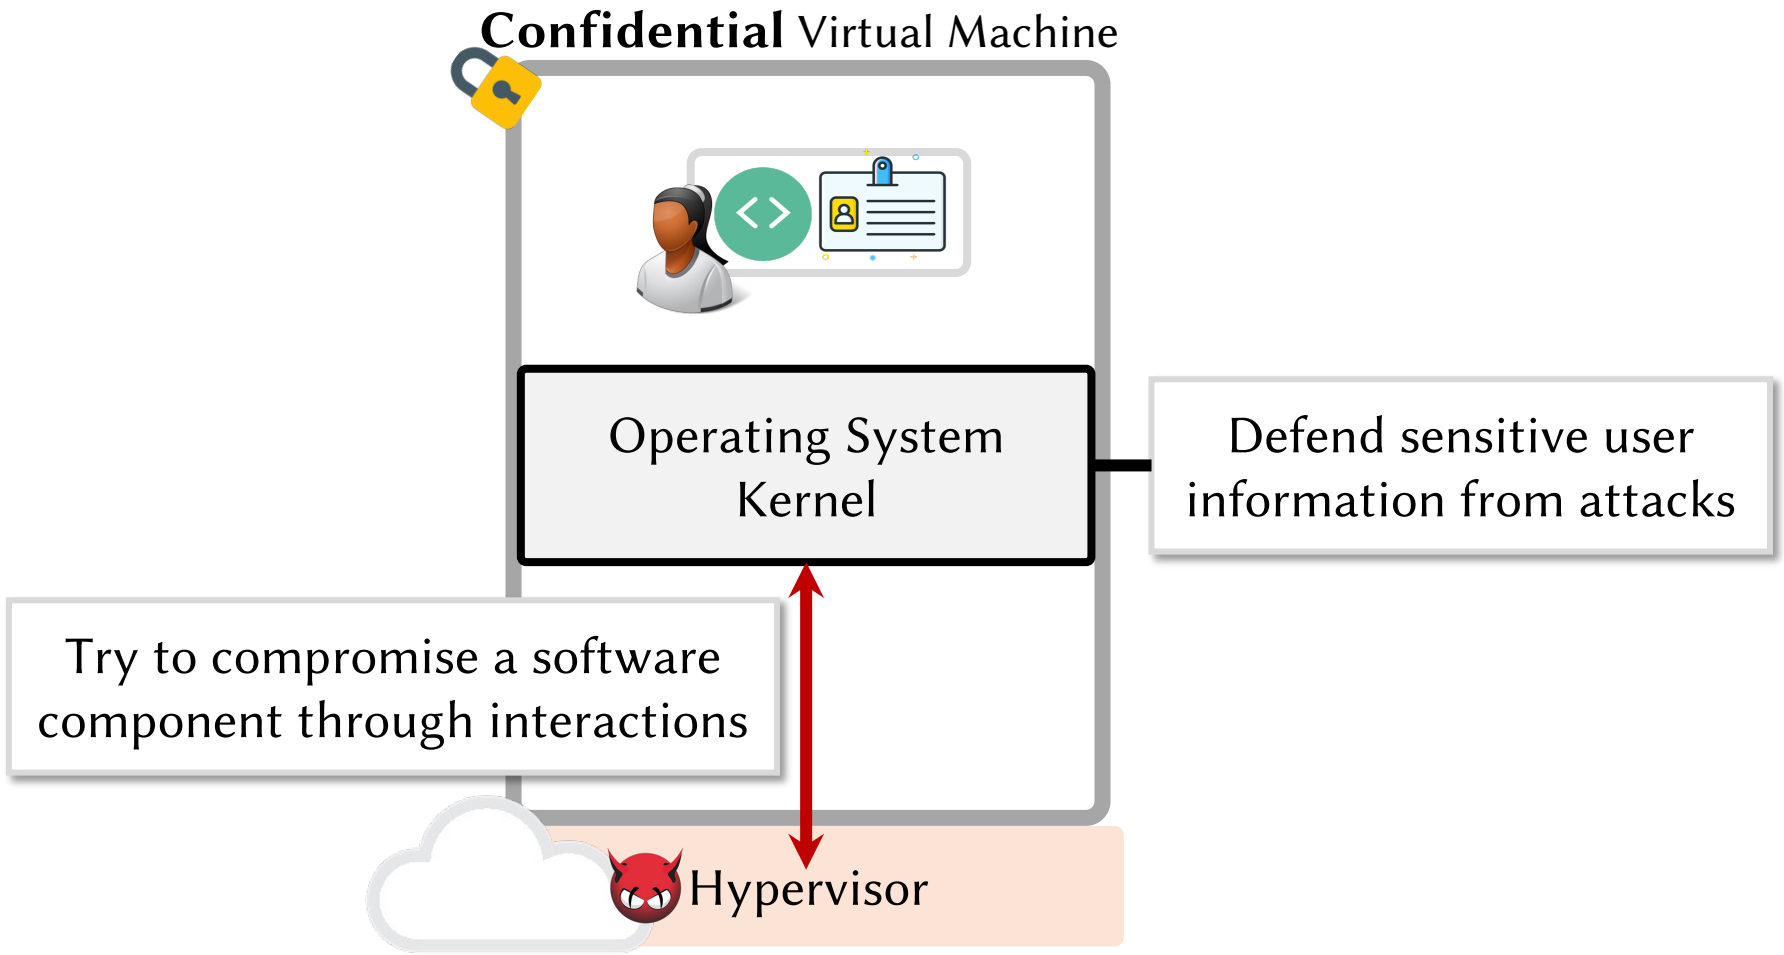

Assumption of trust on the CVM OS is misplaced

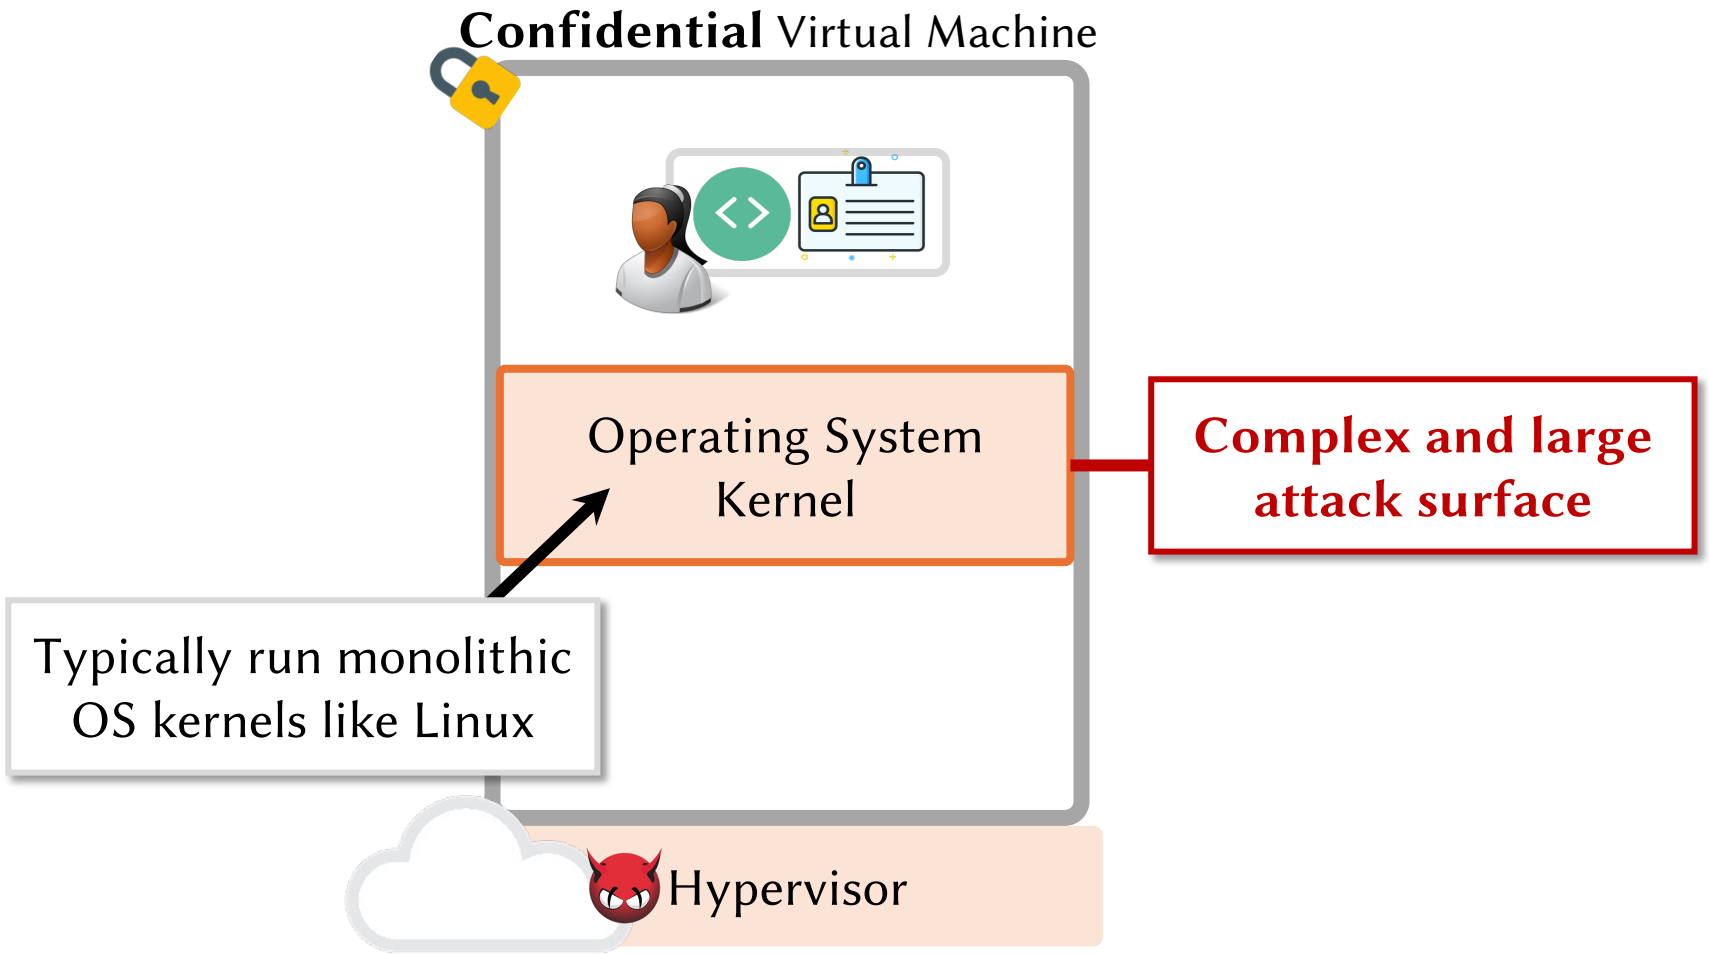

Kernel code integrity as an *example* of misplaced trust

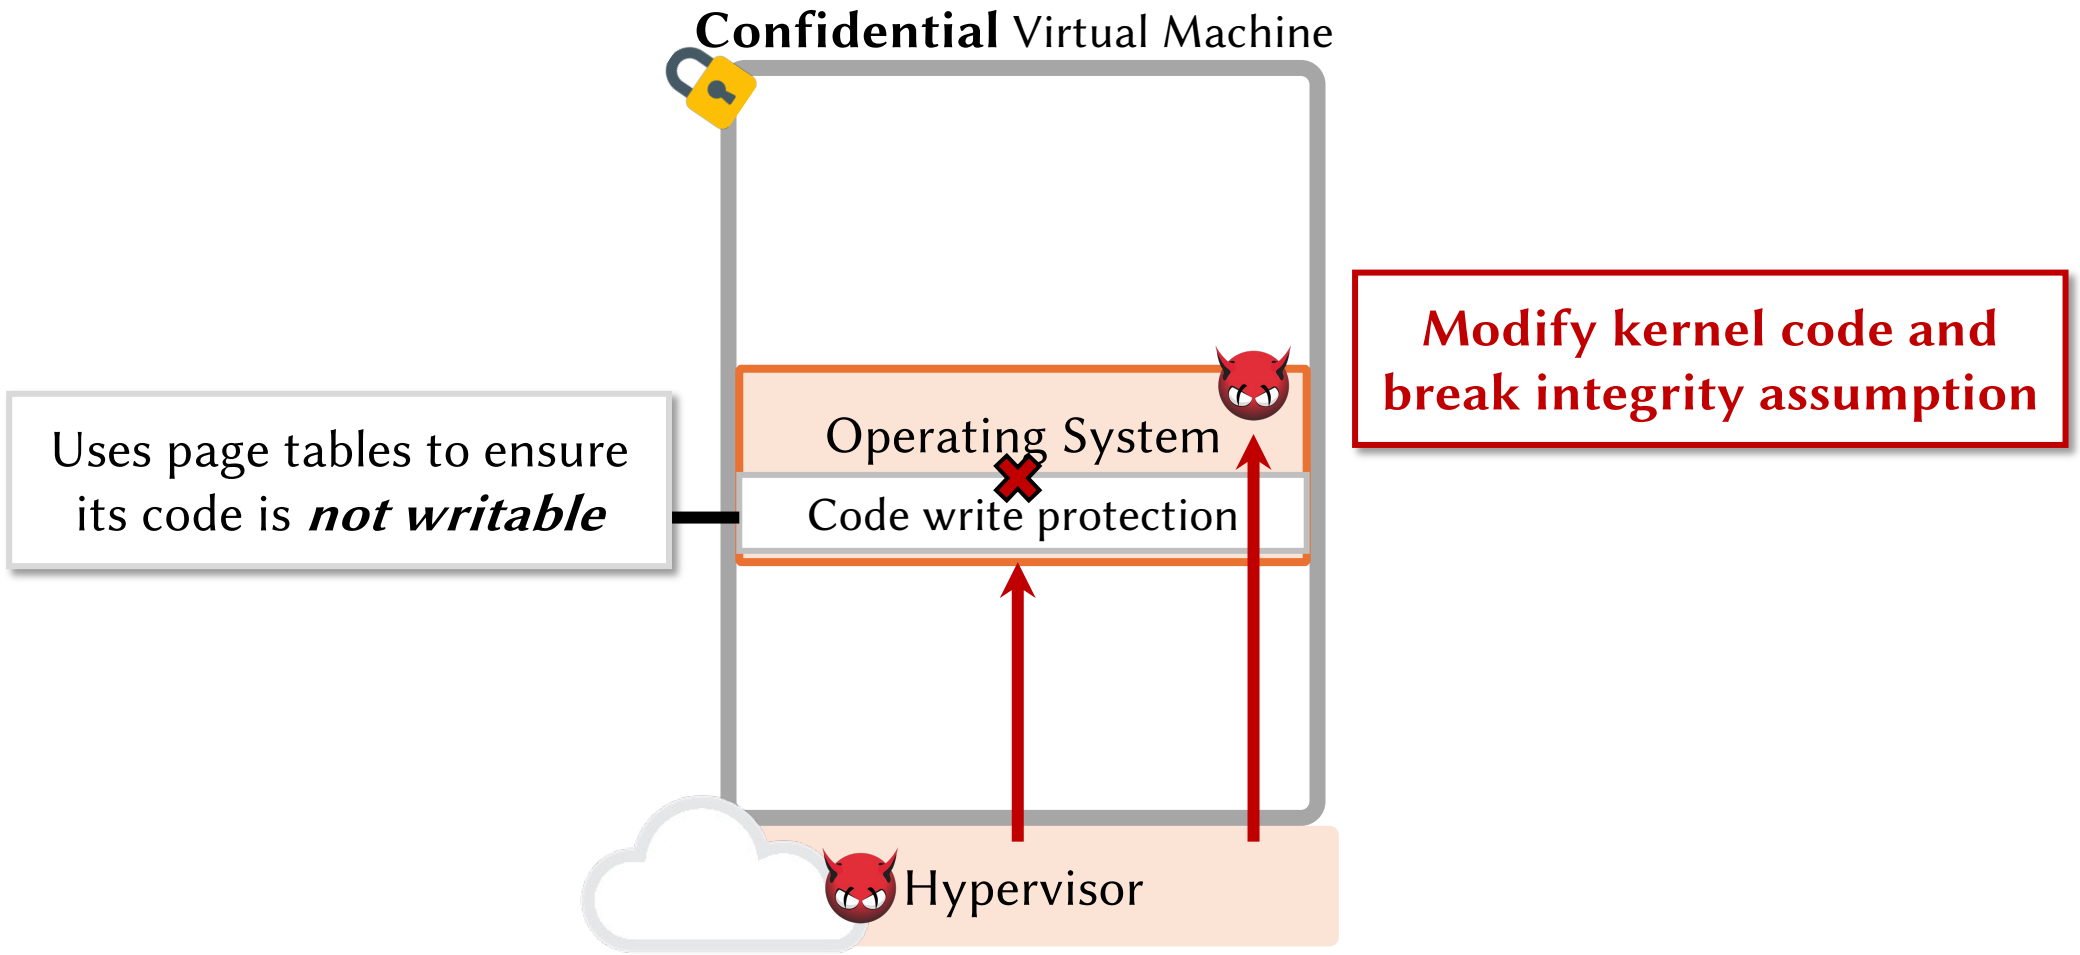

How can we solve the problem of misplaced trust?

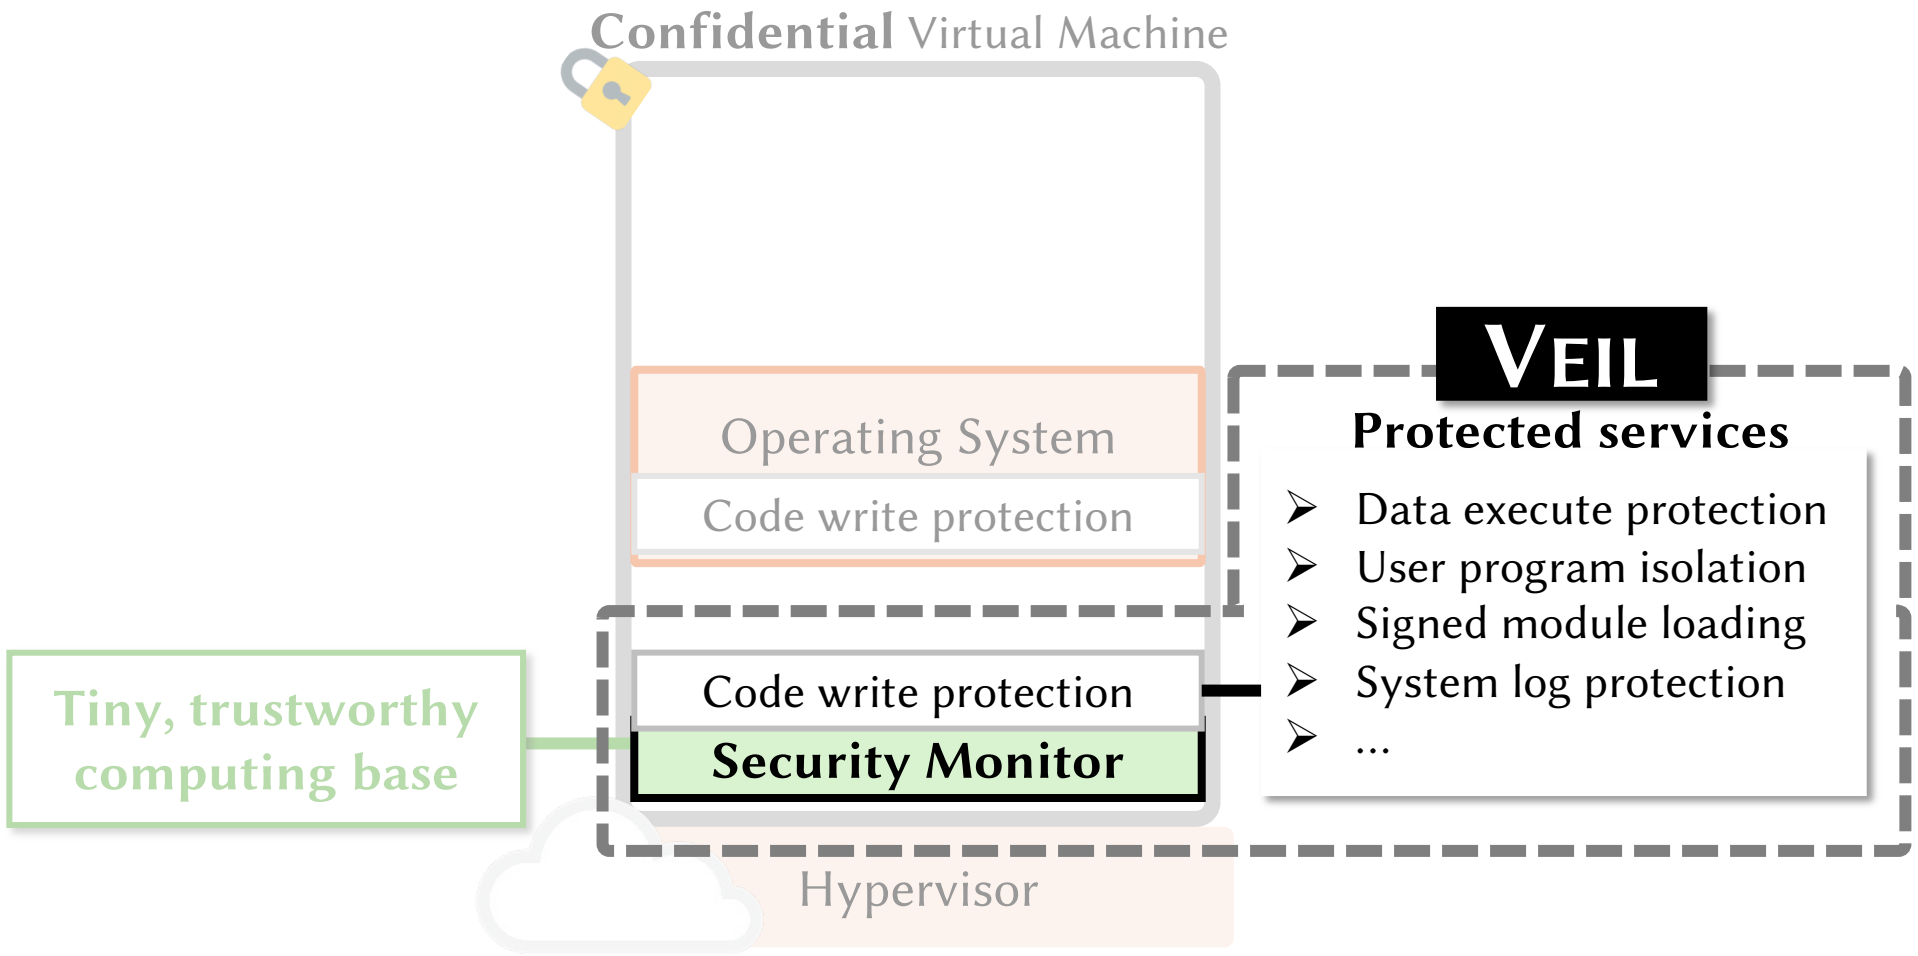

Two requirements for VEIL's monitor and protected services

❑ **Versatile permissions**

➢ Services rely on *read*, *write*, *or execute* restrictions

E.g., code write protection

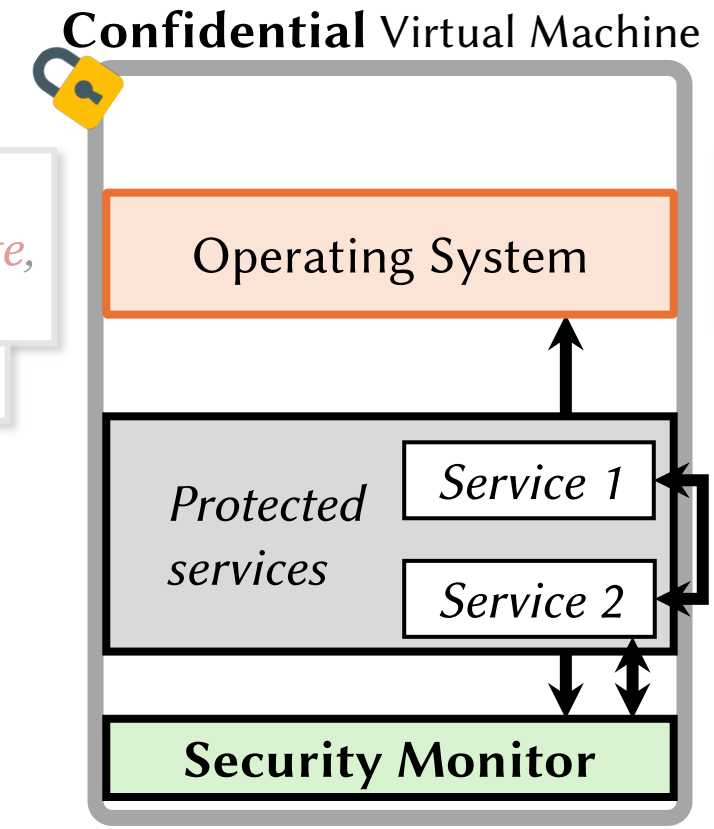

❑ **Multi-domain isolation**

➢ Minimize risk by isolating all components from each other

Isolate

Isolate

Isolate

VEIL leverages Virtual Machine Privilege Levels (VMPL)

- Available in AMD SEV-SNP servers

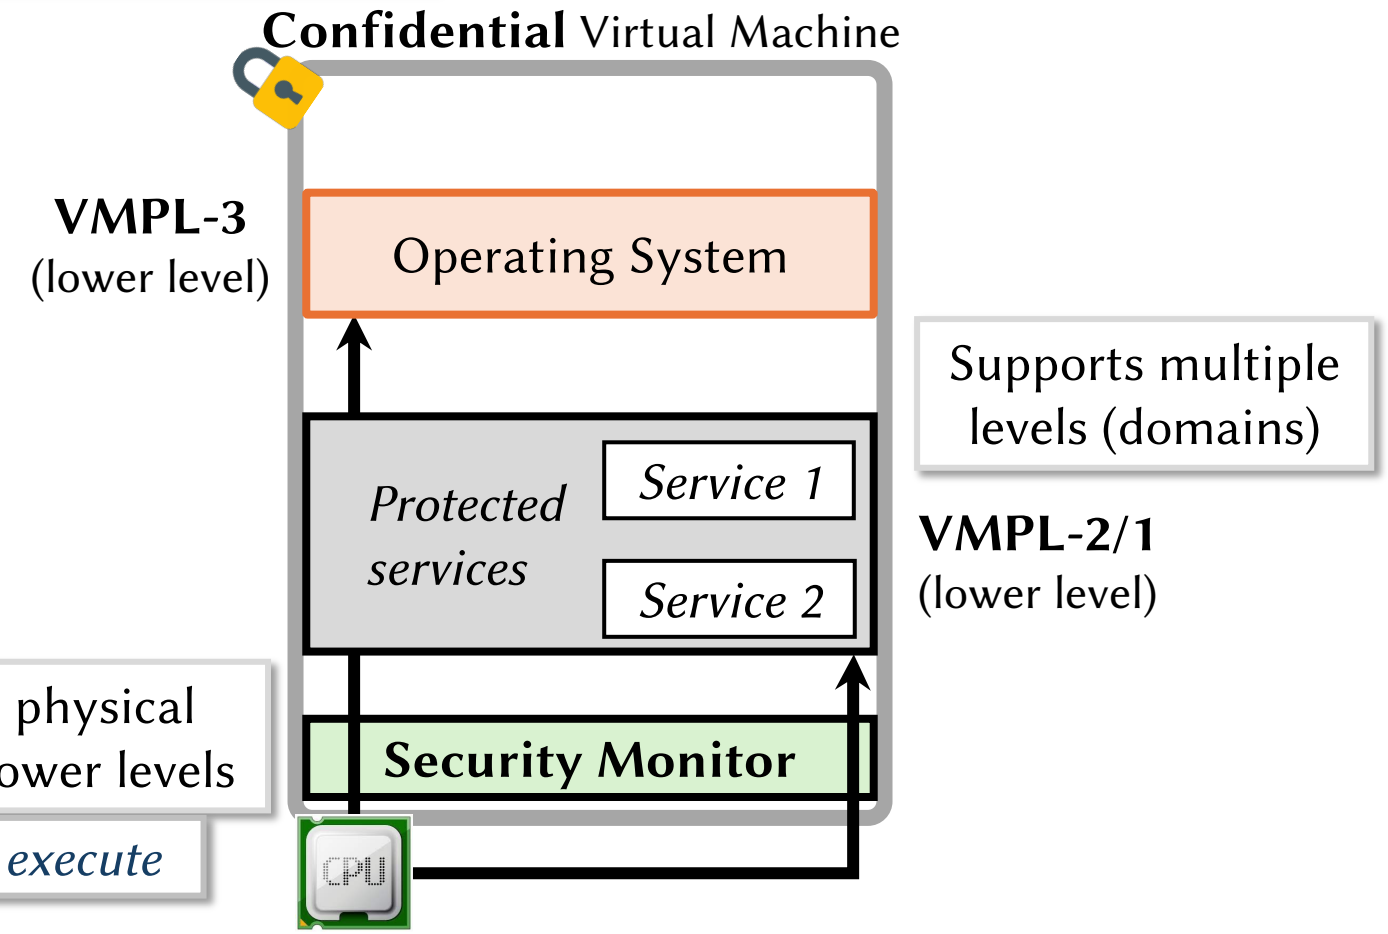

What are the challenges in using VMPL for VEIL?

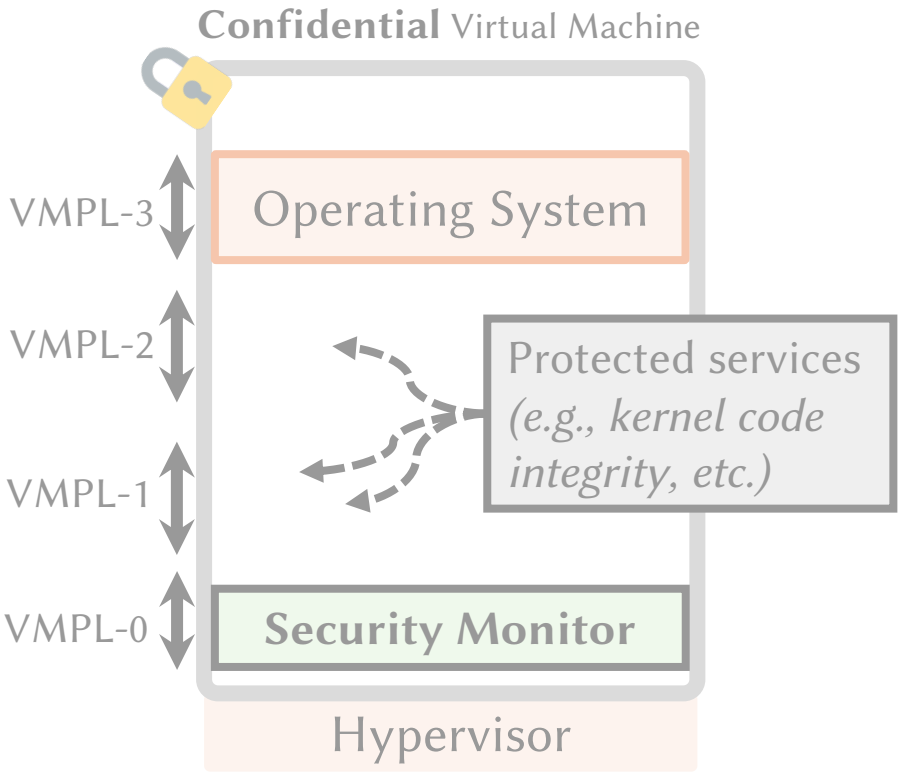

4 domains cannot isolate N>2 protected services

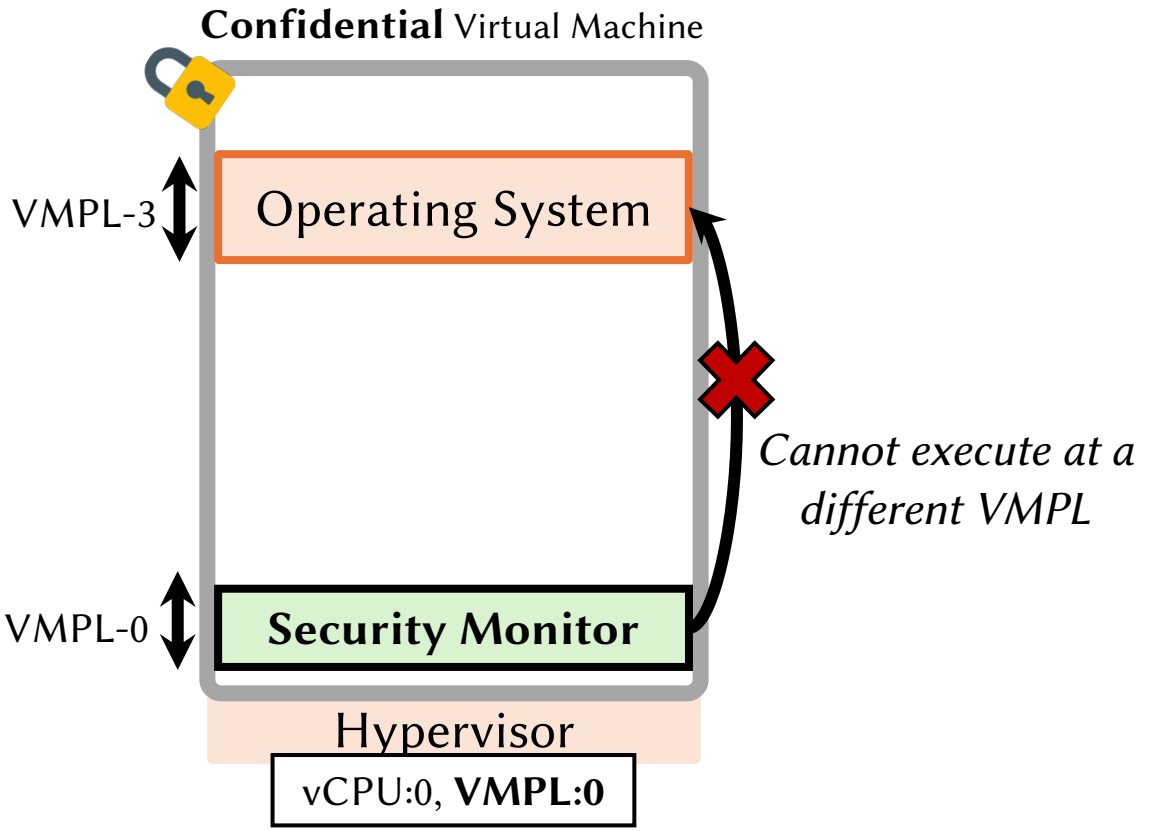

Reserved CPUs at each domain wastes resources

1: Combine VMPL-rings for numerous domains

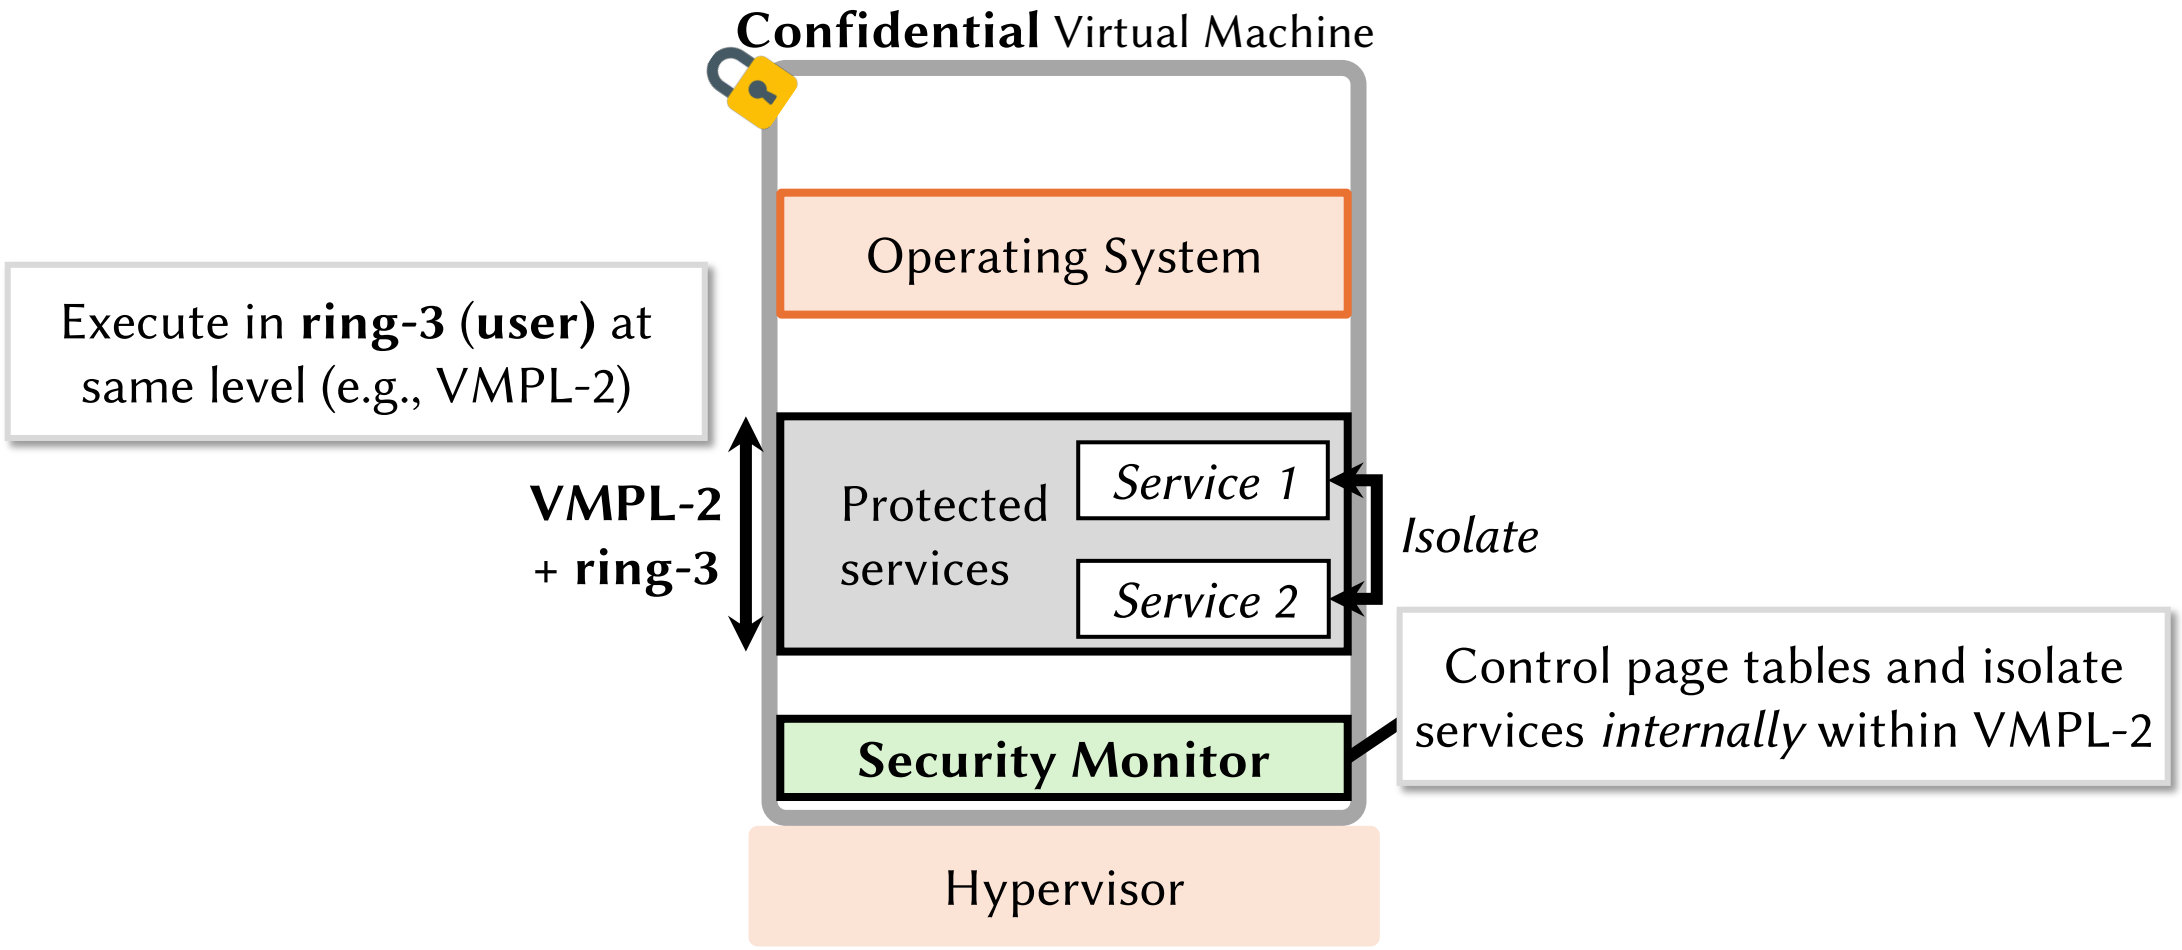

2: Replicate vCPU instances to avoid reservation

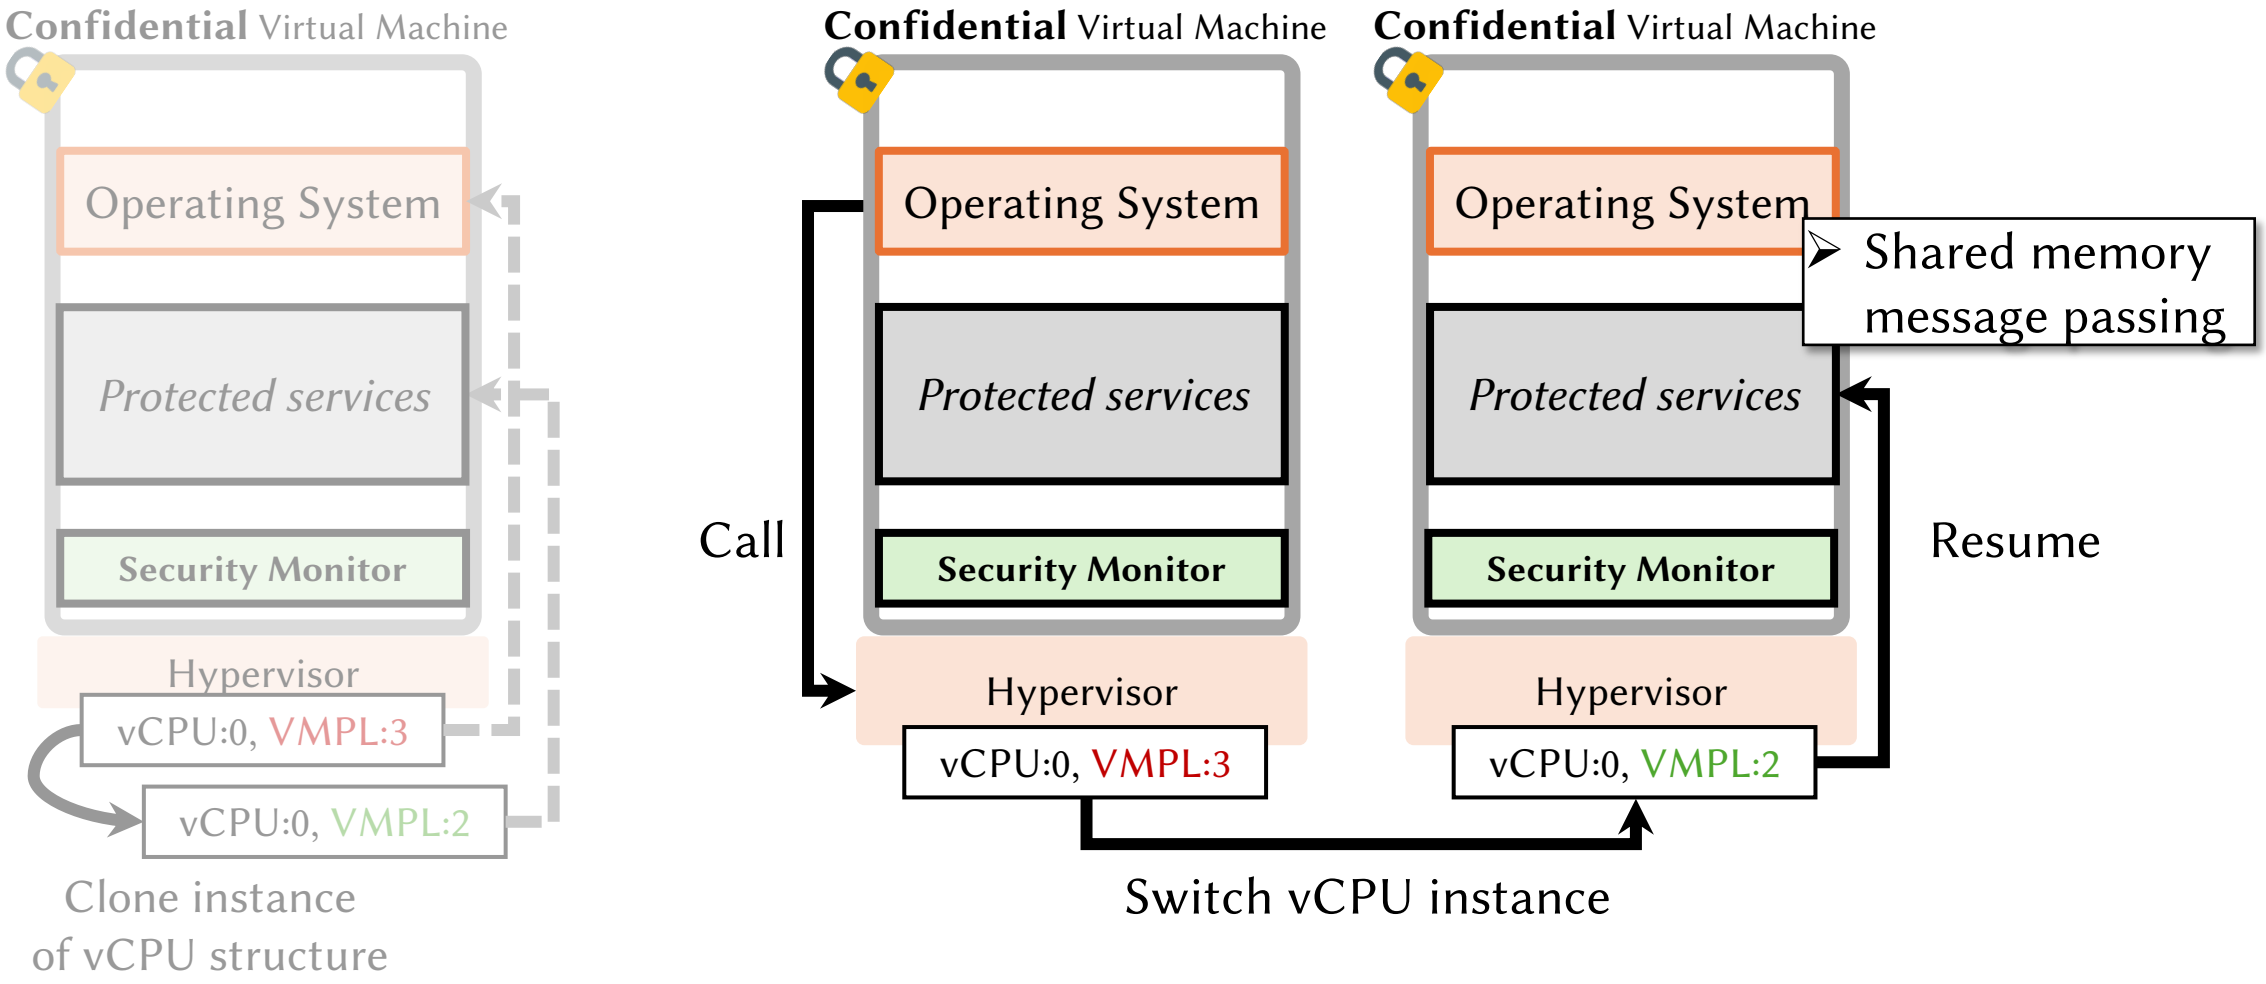

VEIL supports a diverse set of protected services

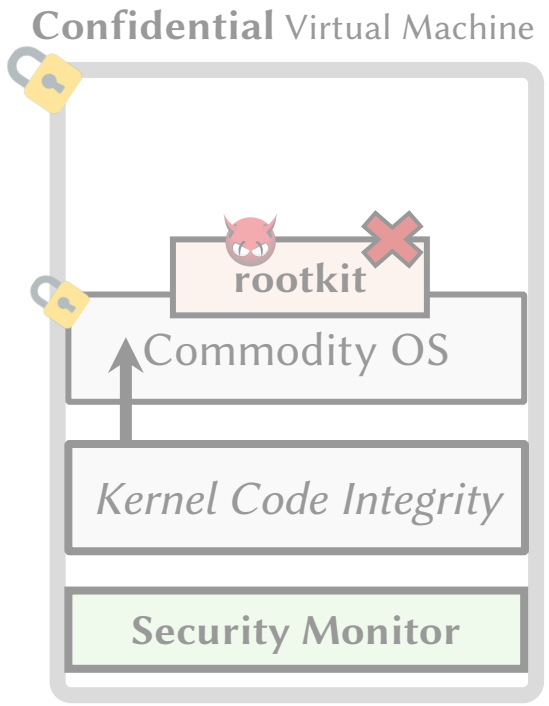

Kernel code integrity
(prevent code injection/rewrite)

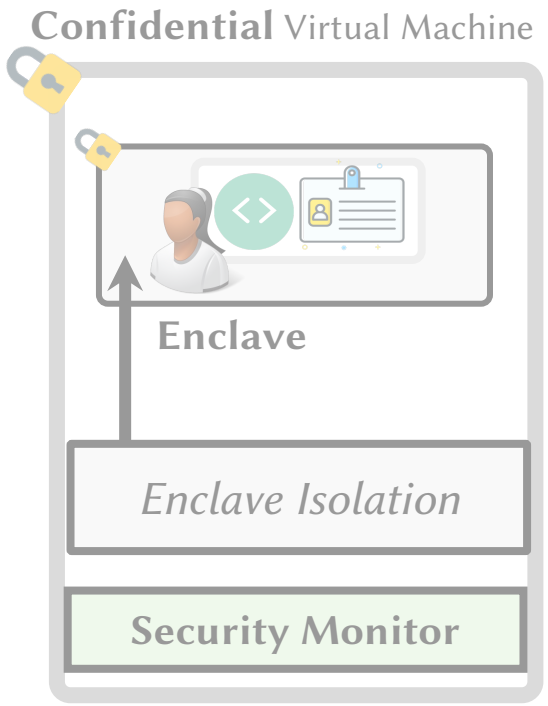

Enclave isolation
(Isolate user programs from OS)

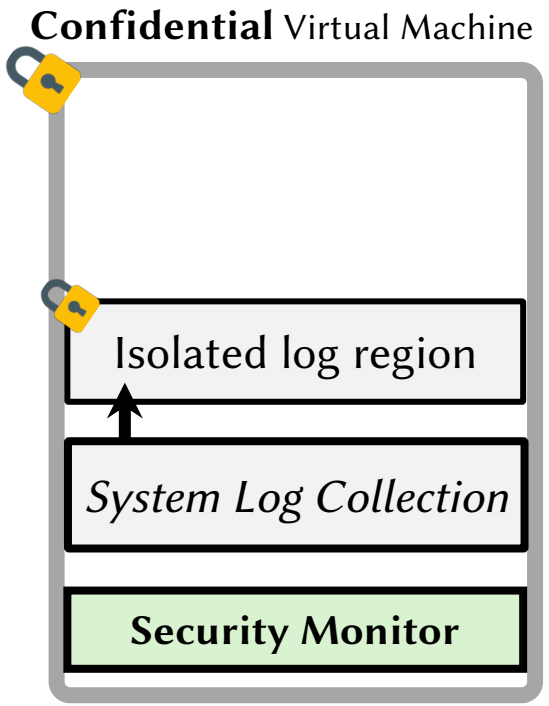

Secure log collection
(store logs in protected regions)

VEIL supports a diverse set of protected services

➢ Will briefly present today

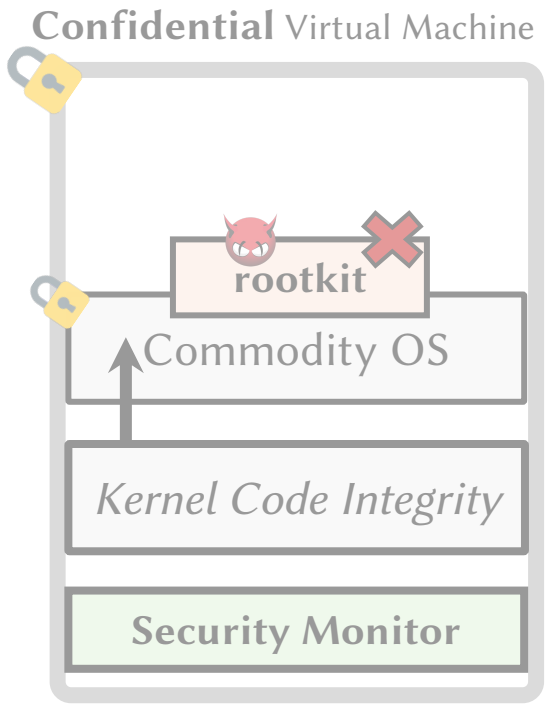

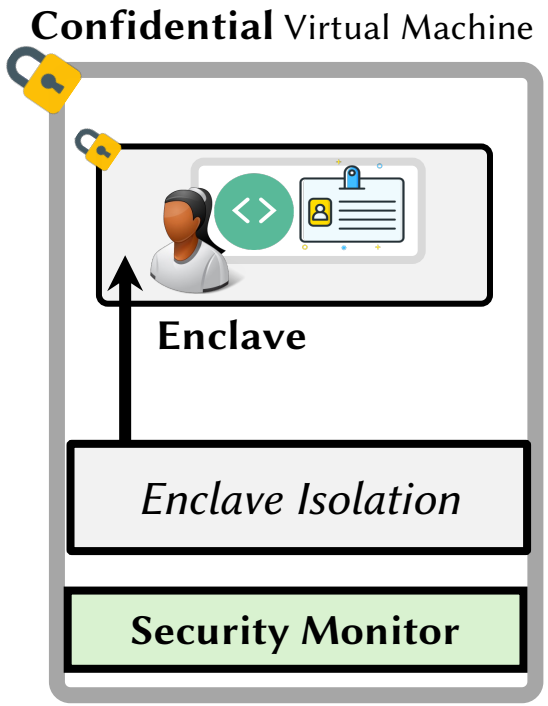

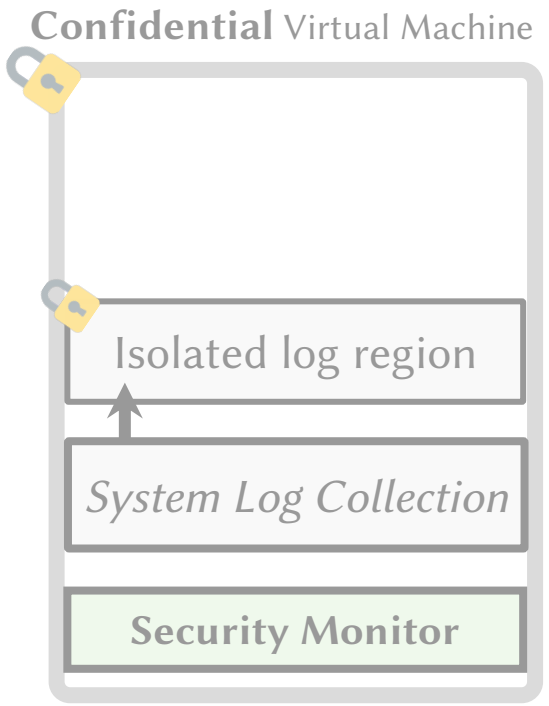

Kernel code integrity
(prevent code injection/rewrite)

Enclave isolation
(Isolate user programs from OS)

Secure log collection:
(store logs in protected regions)

Realize SGX-like enclaves using VEIL in CVMs

Intel Software Guard eXtensions (**SGX**)

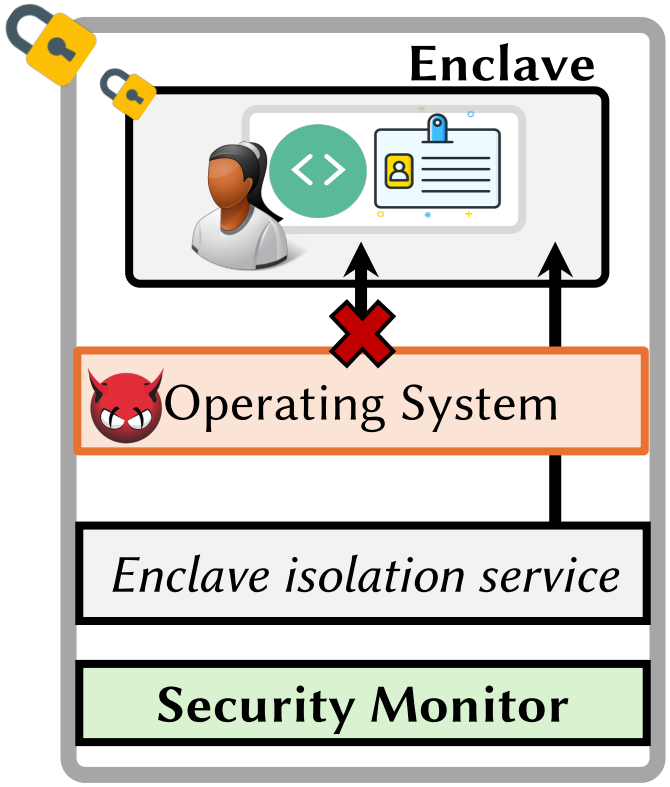

- Memory protect. using VMPL
- Software page table protection
- System call sanitization

What is VEIL's runtime performance overhead?

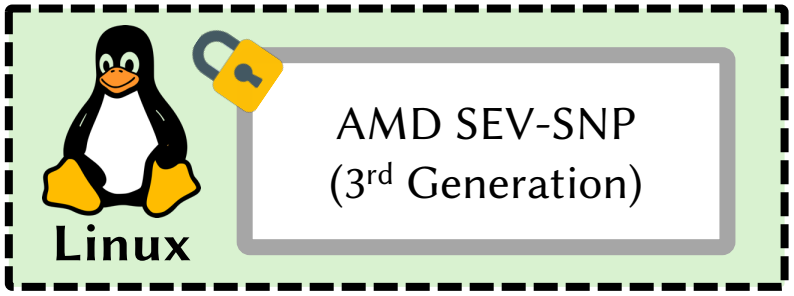

✓ 4 – 6% kernel module loading and unloading increase time

✓ 5 – 69% overhead to isolate programs in enclaves

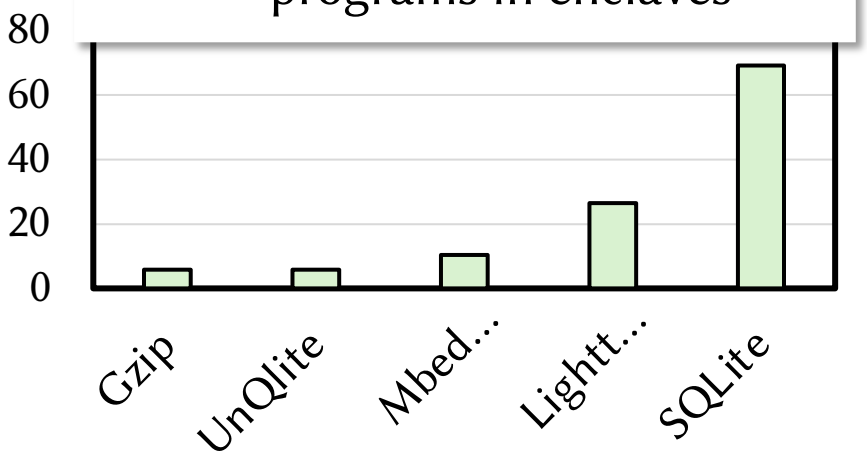

✓ 1 – 19% total overhead to generate and protect logs

➢ Modest overhead for many real-world scenarios

Conclusion

CVMs require a rethink of *trust assignment* for strong security

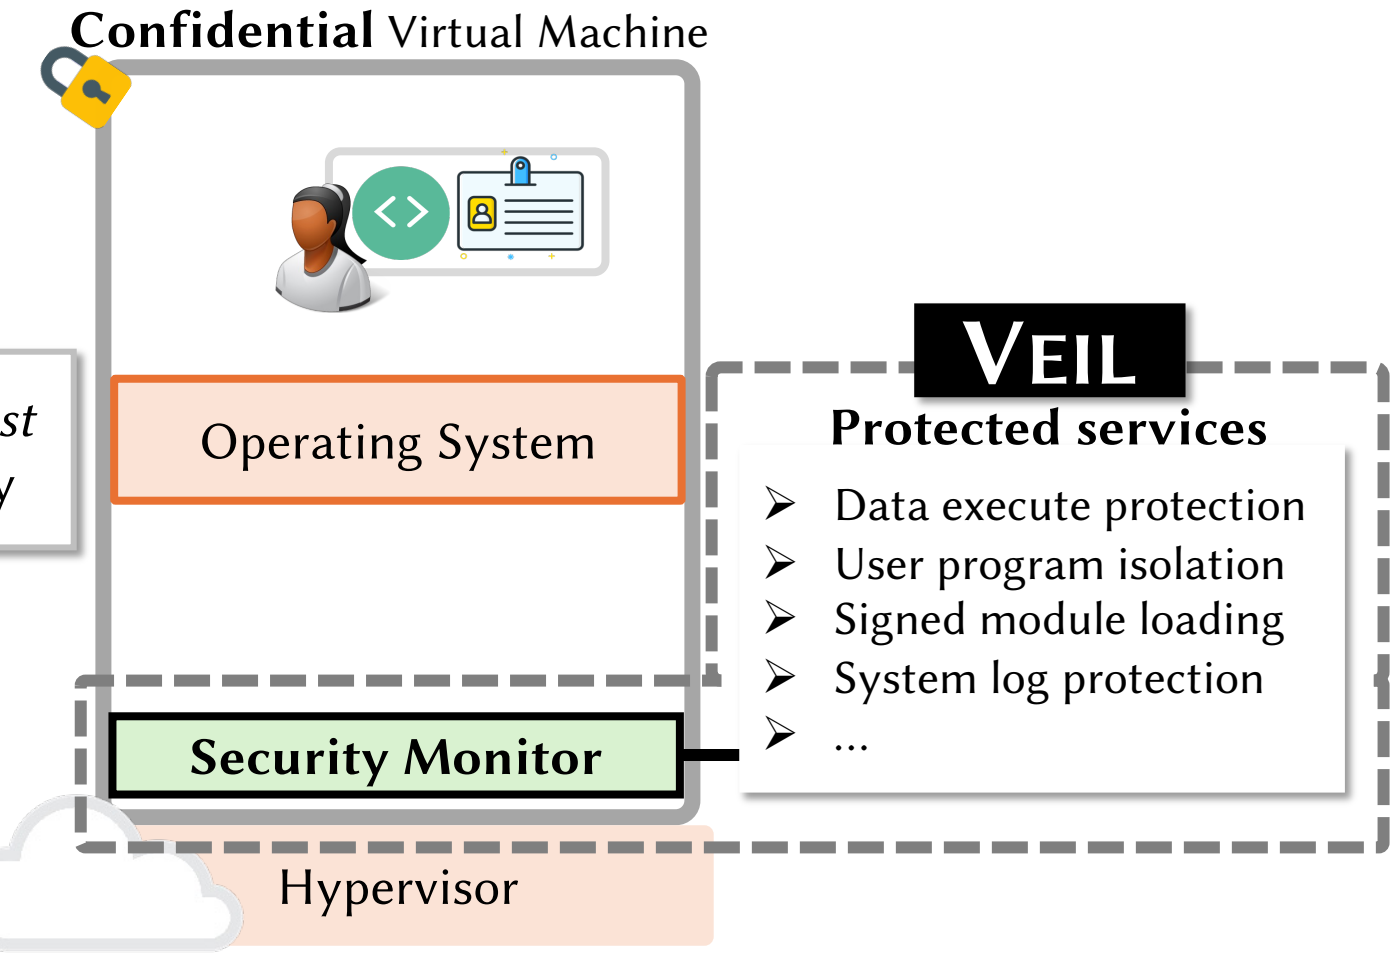

https://github.com/adilahmad17/Veil