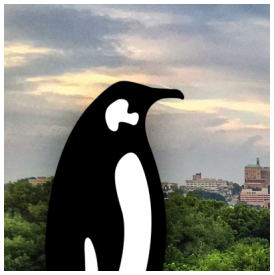

Contribution ID: **171** Type: **not specified**

Hypervisor-Enforced Kernel Integrity (Heki) for KVM

Tuesday, 14 November 2023 09:30 (40 minutes)

Linux kernel vulnerabilities can be mitigated with kernel self-protection mechanisms include control-register pinning and memory page protection restrictions. Unfortunately, none is bullet proof because they are implemented at the same level as the vulnerabilities they try to protect against. To get a more effective defense, we propose to implement some of these protection mechanisms out of the kernel thanks to KVM. Our implementation is inspired by other private implementations currently in use (e.g. Windows's Virtual Secure Mode), but our approach is tailored to Linux specificities.

Taking into account feedback from the first RFC patch series, we are working on a new version bringing minimal static configuration, a new GFN attributes management (alternative to multiple page-track modes), dynamic memory protection for all kernel mappings (e.g. kernel modules, kprobes, ftrace, eBPF JIT), VMM notification, and version management. We'd like to present these new features, discuss the best way to land these changes in mainline, and reach out to potential contributors or users.

Primary author: SALAÜN, Mickaël (Microsoft)

Co-author: Mr VENKATARAMAN, Madhavan (Microsoft)

Presenters: Mr VENKATARAMAN, Madhavan (Microsoft); SALAÜN, Mickaël (Microsoft)

Session Classification: KVM MC

Track Classification: LPC Microconference: KVM MC