Measured Boot, Secure Attestation & co, with systemd

Linux Plumbers Conference 2024, Vienna

Lennart Poettering

Measurements

- Firmware will do some
- systemd-stub will do more
- systemd-pcrextend will do some more in userspace (machine ID, fs IDs, boot phases)

Problem

- Too few PCRs
- All available already assigned uses
- https://uapi-group.org/specifications/specs/linux_tpm_pcr_registry/

Solution: „fake“ PCRs in NV Indexes

- NV Indexes are little slots of (persistent) memory in TPM2, which can be set to various modes
- One such mode is PCR mode: no longer persistent, but can be updated like a PCR.
- In systemd we call such „fake PCRs“ **NvPCRs**
- PCRs are no longer „Beach Front Property“

Allocation

- Dynamic NvPCRs allocation is icky, because we need them early in boot, and need to provide allocation information there.
- Solution: let’s get a static range of NvPCR indexes assigned to systemd from the TCG.
- Actually assigned to UAPI, and further delegated to systemd.
- Static assignments from that range. Easy!

NvPCRs for Everyone

- Not quite (still too few)
- New Measurements: confexts, sysexts, portable services, containers, SMBIOS UUIDs, …

Event Log

- systemd maintains TCG CEL-JSON event log in /run/
- *Almost* TCG CEL, but easily transformable
- systemd-pcrlock builds/validates LUKS unlock policies from this

Event Log

- No rotation right now
- Together with UEFI event log we have fairly complete coverage of what happens on a system
- No clear strategy for kexec, half a stragety for soft-reboot

APIs

- Directly access the event log, it’s API
- Use Varlink IPC call to add additional measurements, „reasonably atomically“.

Next Steps

- Call to generate matching Quote + snapshot of event log
- Tool to verify this (systemd-pcrlock is pretty much this)

Locked Loop

- Quotes travel from node to orchestrator
- systemd-confext DDIs travel from orchestrator to node
- DDIs use dm-verity + dm-crypt and are locked against PCRs just seen, with time window
- Net result: nicely secured configuration deployment, safe for secrets.

That’s all