

Linux Plumbers Conference 2024



Contribution ID: 154

Type: **not specified**

ACPI fast handover for kexec live-update

Wednesday 18 September 2024 15:45 (45 minutes)

It is critical to boot the kexec kernel fast in the system live-update scenarios. As one challenge in this case, a kexec procedure today always initializes ACPI in the same way as a cold reboot, which can take more than 100ms on latest x86 servers. Most of the time is spent on table loading, interpreting and finding idle states, which practically won't have changing side effects across a kexec reboot, as in the case of system software live-update. In order to reduce kexec downtime, an optimization is to preserve (or cache) the ACPI kernel state in memory, eliminating expensive IO and DSDT data structure parsing if the result is guaranteed to be the same.

This requires 1) a mechanism to preserve data between the running kernel and kexec target kernel, (forming an ABI that needs to be carefully defined for compatibility), 2) a new ACPI mode (e.g. a new "acpi=restore" cmdline parameter) that instructs the new kernel to recreate necessary states from memory.

With these extensions, we are looking to further reducing kexec time by 100-300 ms (on top of previous features such as HVO, parallel SMP boot and whitelist based PCI device probing).

This presentation will give an overview of the current state in our ACPI driver, explain the proposed approach at hand including its motivations and constraints, explore and relates to others' work, then discuss possible solutions and look at the next steps.

Primary author: ZHENG, Fam

Presenter: ZHENG, Fam

Session Classification: LPC Refereed Track

Track Classification: LPC Refereed Track