



Contribution ID: 94

Type: **not specified**

Kernel handling of CPU and memory hot un/plug events for crash

Wednesday, 15 November 2023 09:30 (45 minutes)

Once the kdump service is loaded, if changes to CPUs or memory occur, either by hot un/plug or off/onlining, the crash elfcorehdr must also be updated.

The elfcorehdr describes to kdump the CPUs and memory in the system, and any inaccuracies can result in a vmcore with missing CPU context or memory regions.

The current solution utilizes a udev event per CPU or memblock to initiate an unload-then-reload of the kdump image (eg. kernel, initrd, boot_params, purgatory and elfcorehdr) by the userspace kexec utility. In a rapidly scaling environment, significant performance problems occur related to offloading this activity to userspace.

This talk introduces a generic crash handler that registers with the CPU and memory notifiers. Upon CPU or memory changes, from either hot un/plug or off/onlining, this generic handler is invoked and performs important housekeeping, for example obtaining the appropriate lock, and then invokes an architecture specific handler to do the appropriate elfcorehdr update.

Primary authors: DEVOLDER, Eric (Oracle); JAIN, Sourabh

Presenters: DEVOLDER, Eric (Oracle); JAIN, Sourabh

Session Classification: Kernel Summit

Track Classification: Kernel Summit Track