



Contribution ID: 118

Type: **not specified**

Use of Netlink for thermal kernel-user notification is problematic

Tuesday, 14 November 2023 15:10 (20 minutes)

Use of Netlink for thermal kernel-user notification is problematic

Netlink is used in thermal subsystem to notify events. This includes trip changes, temperature change, capability changes etc.

There are two notify protocols used:

NETLINK_GENERIC : Added couple of years back as the preferred kernel-user notification mechanism

NETLINK_KOBJECT_UEVENT: Used by thermal user space governor

Both are problematic. It is assumed that kernel will not have notification storm. On some x86 there can be 10s of thermal zones. If there is a storm of events, the system freezes. Only way to recover is force reboot. This was observed in some customer system, where firmware sent too many events waiting for user space to act.

NETLINK_KOBJECT_UEVENT creates more problem. The udev daemon creates many workers to handle a storm, which is related to number of CPUs. This increases dynamic memory foot print.

In my experiments if netlink events storm with a period of a milli second or less, system totally freezes.

Primary author: PANDRUVADA, Srinivas

Presenter: PANDRUVADA, Srinivas

Session Classification: Power Management and Thermal Control MC

Track Classification: LPC Microconference: Power Management and Thermal Control MC