



Contribution ID: 208

Type: **not specified**

Workqueues and CPU Hotplug

Wednesday, 14 November 2018 11:30 (15 minutes)

Flexible workqueue: Currently we have two pool setting-up for workqueue: 1) per-cpu workqueue pool and 2) unbound workqueue pool, the former require the users of workqueues to have some knowledge of cpu online state, as shown in:

<https://lore.kernel.org/lkml/20180625224332.10596-2-paulmck@linux.vnet.ibm.com/T/#u>

While the latter (unbound workqueue) only has one pool per-NUMA, and that may hurt the scalability if we want to run multiple tasks in parallel inside a NUMA node.

Therefore, that is a clear requirement for having a setting-up for workqueue to provide flexible level of parallelism (i.e. could run as many tasks as possible while save users from worrying about race with cpu hotplug).

We'd like to have a session to talk about requirement and possible solution.

I agree to abide by the anti-harassment policy

Presenter: FENG, Boqun

Session Classification: Performance and Scalability MC