
Subject: Re: v2.6.26-rc7/cgroups: circular locking dependency
Posted by [Peter Zijlstra](#) on Sun, 22 Jun 2008 15:50:21 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Mon, 2008-06-23 at 00:34 +0900, KOSAKI Motohiro wrote:

> CC'ed Paul Jackson

>

> it seems typical ABBA deadlock.

> I think cpuset use cgrou_lock() by mistake.

>

> IMHO, cpuset_handle_cpuhp() sholdn't use cgroup_lock() and

> shouldn't call rebuild_sched_domains().

Looks like Max forgot to test with lockdep enabled...

Well, someone should when you change the online map.

Max, Paul, can we handle this in update_sched_domains() instead?

```
> -> #1 (cgroup_mutex){--}:
>      [<c015a435>] __lock_acquire+0xf45/0x1040
>      [<c015a5c8>] lock_acquire+0x98/0xd0
>      [<c05416d1>] mutex_lock_nested+0xb1/0x300
>      [<c0160e6f>] cgroup_lock+0xf/0x20    cgroup_lock
>      [<c0164750>] cpuset_handle_cpuhp+0x20/0x180
>      [<c014ea77>] notifier_call_chain+0x37/0x70
>      [<c014eae9>] __raw_notifier_call_chain+0x19/0x20
>      [<c051f8c8>] _cpu_down+0x78/0x240  cpu_hotplug.lock
>      [<c051fabb>] cpu_down+0x2b/0x40    cpu_add_remove_lock
>      [<c0520cd9>] store_online+0x39/0x80
>      [<c02f627b>] sysdev_store+0x2b/0x40
>      [<c01d3372>] sysfs_write_file+0xa2/0x100
>      [<c0195486>] vfs_write+0x96/0x130
>      [<c0195b4d>] sys_write+0x3d/0x70
>      [<c010831b>] sysenter_past_esp+0x78/0xd1
>      [<ffffffff>] 0xffffffff
>
> -> #0 (&cpu_hotplug.lock){--}:
>      [<c0159fe5>] __lock_acquire+0xaf5/0x1040
>      [<c015a5c8>] lock_acquire+0x98/0xd0
>      [<c05416d1>] mutex_lock_nested+0xb1/0x300
>      [<c015efbc>] get_online_cpus+0x2c/0x40    cpu_hotplug.lock
>      [<c0163e6d>] rebuild_sched_domains+0x7d/0x3a0
>      [<c01653a4>] cpuset_common_file_write+0x204/0x440 cgroup_lock
>      [<c0162bc7>] cgroup_file_write+0x67/0x130
>      [<c0195486>] vfs_write+0x96/0x130
>      [<c0195b4d>] sys_write+0x3d/0x70
>      [<c010831b>] sysenter_past_esp+0x78/0xd1
```

```
>      [<ffffffff>] 0xffffffff
>
>
> > Hi,
> >
> > I decided to see what cgroups is all about, and followed the instructions
> > in Documentation/cgroups.txt :-) It happened when I did this:
> >
> > [root@damson /dev/cgroup/Vegard 0]
> > # echo 1 > cpuset.cpus
> >
> > I can also provide the kernel config if necessary.
> >
> >
> > Vegard
> >
> >
> > =====
> > [ INFO: possible circular locking dependency detected ]
> > 2.6.26-rc7 #25
> > -----
> > bash/10032 is trying to acquire lock:
> > (&cpu_hotplug.lock){--...}, at: [<c015efbc>] get_online_cpus+0x2c/0x40
> >
> > but task is already holding lock:
> > (cgroup_mutex){--...}, at: [<c0160e6f>] cgroup_lock+0xf/0x20
> >
> > which lock already depends on the new lock.
> --
> To unsubscribe from this list: send the line "unsubscribe linux-kernel" in
> the body of a message to majordomo@vger.kernel.org
> More majordomo info at http://vger.kernel.org/majordomo-info.html
> Please read the FAQ at http://www.tux.org/lkml/
```

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
