

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: [
```

```
but task is already holding lock:
(cgroup_mutex){--..}, at: [
```

which lock already depends on the new lock.

the existing dependency chain (in reverse order) is:

```
-> #1 (cgroup_mutex){--..}:
  [
```

```
[<c010831b>] sysenter_past_esp+0x78/0xd1  
[<ffffff>] 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  
[<c0163e6d>] rebuild_sched_domains+0x7d/0x3a0  
[<c01653a4>] cpuset_common_file_write+0x204/0x440  
[<c0162bc7>] cgroup_file_write+0x67/0x130  
[<c0195486>] vfs_write+0x96/0x130  
[<c0195b4d>] sys_write+0x3d/0x70  
[<c010831b>] sysenter_past_esp+0x78/0xd1  
[<ffffff>] 0xffffffff
```

other info that might help us debug this:

1 lock held by bash/10032:

```
#0: (cgroup_mutex){--..}, at: [<c0160e6f>] cgroup_lock+0xf/0x20
```

stack backtrace:

```
Pid: 10032, comm: bash Not tainted 2.6.26-rc7 #25  
[<c0157ae7>] print_circular_bug_tail+0x77/0x90  
[<c0159fe5>] __lock_acquire+0xaf5/0x1040  
[<c015a5c8>] lock_acquire+0x98/0xd0  
[<c015efbc>] ? get_online_cpus+0x2c/0x40  
[<c05416d1>] mutex_lock_nested+0xb1/0x300  
[<c015efbc>] ? get_online_cpus+0x2c/0x40  
[<c015efbc>] ? get_online_cpus+0x2c/0x40  
[<c015efbc>] get_online_cpus+0x2c/0x40  
[<c0163e6d>] rebuild_sched_domains+0x7d/0x3a0  
[<c0158e65>] ? mark_held_locks+0x65/0x80  
[<c0190ac5>] ? kfree+0xb5/0x120  
[<c027736a>] ? heap_free+0xa/0x10  
[<c027736a>] ? heap_free+0xa/0x10  
[<c01653a4>] cpuset_common_file_write+0x204/0x440  
[<c0163c30>] ? started_after+0x0/0x70  
[<c010e585>] ? native_sched_clock+0xb5/0x110  
[<c0162e10>] ? started_after+0x0/0x70  
[<c0164c80>] ? cpuset_test_cpumask+0x0/0x20  
[<c0164ca0>] ? cpuset_change_cpumask+0x0/0x20  
[<c0162bc7>] cgroup_file_write+0x67/0x130  
[<c01651a0>] ? cpuset_common_file_write+0x0/0x440  
[<c010e585>] ? native_sched_clock+0xb5/0x110  
[<c0195486>] vfs_write+0x96/0x130  
[<c0162b60>] ? cgroup_file_write+0x0/0x130  
[<c0195b4d>] sys_write+0x3d/0x70
```

[<c010831b>] sysenter_past_esp+0x78/0xd1

=====

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

Subject: Re: v2.6.26-rc7/cgroups: circular locking dependency
Posted by [Cyrill Gorcunov](#) on Sun, 22 Jun 2008 09:10:13 GMT
[View Forum Message](#) <> [Reply to Message](#)

[Vegard Nossum - Sat, Jun 21, 2008 at 07:38:59PM +0200]

| 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.

| the existing dependency chain (in reverse order) is:

```
| -> #1 (cgroup_mutex){--..}:  
|   [<c015a435>] __lock_acquire+0xf45/0x1040  
|   [<c015a5c8>] lock_acquire+0x98/0xd0  
|   [<c05416d1>] mutex_lock_nested+0xb1/0x300  
|   [<c0160e6f>] cgroup_lock+0xf/0x20
```

```
[<c0164750>] cpuset_handle_cpuhp+0x20/0x180
[<c014ea77>] notifier_call_chain+0x37/0x70
[<c014eae9>] __raw_notifier_call_chain+0x19/0x20
[<c051f8c8>] _cpu_down+0x78/0x240
[<c051fabb>] cpu_down+0x2b/0x40
[<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
[<c0163e6d>] rebuild_sched_domains+0x7d/0x3a0
[<c01653a4>] cpuset_common_file_write+0x204/0x440
[<c0162bc7>] cgroup_file_write+0x67/0x130
[<c0195486>] vfs_write+0x96/0x130
[<c0195b4d>] sys_write+0x3d/0x70
[<c010831b>] sysenter_past_esp+0x78/0xd1
[<ffffffff>] 0xffffffff
```

other info that might help us debug this:

1 lock held by bash/10032:

#0: (cgroup_mutex){--..}, at: [<c0160e6f>] cgroup_lock+0xf/0x20

stack backtrace:

Pid: 10032, comm: bash Not tainted 2.6.26-rc7 #25

```
[<c0157ae7>] print_circular_bug_tail+0x77/0x90
[<c0159fe5>] __lock_acquire+0xaf5/0x1040
[<c015a5c8>] lock_acquire+0x98/0xd0
[<c015efbc>] ? get_online_cpus+0x2c/0x40
[<c05416d1>] mutex_lock_nested+0xb1/0x300
[<c015efbc>] ? get_online_cpus+0x2c/0x40
[<c015efbc>] ? get_online_cpus+0x2c/0x40
[<c015efbc>] get_online_cpus+0x2c/0x40
[<c0163e6d>] rebuild_sched_domains+0x7d/0x3a0
[<c0158e65>] ? mark_held_locks+0x65/0x80
[<c0190ac5>] ? kfree+0xb5/0x120
[<c027736a>] ? heap_free+0xa/0x10
[<c027736a>] ? heap_free+0xa/0x10
[<c01653a4>] cpuset_common_file_write+0x204/0x440
[<c0163c30>] ? started_after+0x0/0x70
```

```
| [<c010e585>] ? native_sched_clock+0xb5/0x110
| [<c0162e10>] ? started_after+0x0/0x70
| [<c0164c80>] ? cpuset_test_cpumask+0x0/0x20
| [<c0164ca0>] ? cpuset_change_cpumask+0x0/0x20
| [<c0162bc7>] cgroup_file_write+0x67/0x130
| [<c01651a0>] ? cpuset_common_file_write+0x0/0x440
| [<c010e585>] ? native_sched_clock+0xb5/0x110
| [<c0195486>] vfs_write+0x96/0x130
| [<c0162b60>] ? cgroup_file_write+0x0/0x130
| [<c0195b4d>] sys_write+0x3d/0x70
| [<c010831b>] sysenter_past_esp+0x78/0xd1
| =====
|
```

Hi Vegard,

would be interesting to see your config please

- Cyrill -

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

Subject: Re: v2.6.26-rc7/cgroups: circular locking dependency
Posted by [Vegard Nossum](#) on Sun, 22 Jun 2008 09:42:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Sun, Jun 22, 2008 at 11:10 AM, Cyrill Gorcunov <gorcunov@gmail.com> wrote:
> [Vegard Nossum - Sat, Jun 21, 2008 at 07:38:59PM +0200]

```
> | =====
> | [ 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
```

...

```
> Hi Vegard,
>
> would be interesting to see your config please
```

Here is my config and vmlinux:

<http://www.kernel.org/pub/linux/kernel/people/vegard/kernels/20080621-v2.6.26-rc7/>

By the way, I was doing CPU hotplug/unplug at the time.

Thanks :)

Vegard

--

"The animistic metaphor of the bug that maliciously sneaked in while the programmer was not looking is intellectually dishonest as it disguises that the error is the programmer's own creation."

-- E. W. Dijkstra, EWD1036

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

Subject: Re: v2.6.26-rc7/cgroups: circular locking dependency
Posted by [Cyrill Gorcunov](#) on Sun, 22 Jun 2008 09:50:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

[Vegard Nossum - Sun, Jun 22, 2008 at 11:42:50AM +0200]
| On Sun, Jun 22, 2008 at 11:10 AM, Cyrill Gorcunov <gorcunov@gmail.com> wrote:
| > [Vegard Nossum - Sat, Jun 21, 2008 at 07:38:59PM +0200]
| > | =====
| > | [INFO: possible circular locking dependency detected]
| > | 2.6.26-rc7 #25
| > | -----
| > | bash/10032 is trying to acquire lock:
| > | (&cpu_hotplug.lock){--..}, at: [| > |
| > | but task is already holding lock:
| > | (cgroup_mutex){--..}, at: [| > |
| ...
| > Hi Vegard,
| >
| > would be interesting to see your config please
| >
| Here is my config and vmlinux:
| <http://www.kernel.org/pub/linux/kernel/people/vegard/kernels/20080621-v2.6.26-rc7/>
|

| By the way, I was doing CPU hotplug/unplug at the time.

| Thanks :)

| Vegard

| --

| "The animistic metaphor of the bug that maliciously sneaked in while
| the programmer was not looking is intellectually dishonest as it
| disguises that the error is the programmer's own creation."

| -- E. W. Dijkstra, EWD1036

Thanks Vegard,

actually i'm not that spec in this area ;) but will take a look

- Cyrill -

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
