
Subject: Re: PPC64 Kernel 2.6.18 + RPM packages

Posted by [dev](#) on Fri, 25 May 2007 14:04:35 GMT

[View Forum Message](#) <> [Reply to Message](#)

Christian, please check the patch attached.

Kirill

Christian Kaiser2 wrote:

> Kirill Korotaev <dev@sw.ru> wrote on 05/25/2007 02:41:09 PM:

>

>

>>Christian,

>>

>>sys_fairsched_chwt() can return ENOENT in the only case:

>>when fairsched node ID was not found in the list.

>>but VE0 node is always on the list.

>>So I guess there is something wrong with syscall args in your case.

>>

>>Please apply the attached patch to the kernel and vzctl

>>recompile/reboot and report the output of both.

>>

>>Thanks,

>>Kirill

>>

>>

>

>

> OK! here it is:

>

> 1. vzctl:

>

> [root@s10b10-0 ~]# vzctl set 0 --cpuunits 1000

> Setting CPU units: 1000

> vzctl: fairsched_chwt(2147483647, 500)

> fairsched_chwt: No such file or directory

>

> 2. kernel messages (dmesg):

>

> fairsched_find: search 500, node: 2147483647/2147483647

>

>

>>>When, where and what should return "-ENOSYS"?

>>>

>>>I rechecked my .config file and I've seen that

>>>CONFIG_SCHED_VCPU=y

>>>CONFIG_FAIRSCHEDED=y

>>>is already in it. But I cannot find it in the menuconfig. The menu of

```
>>>powerpc architecture is a bit different than i386. So perhaps I just
>>>haven't seen it... where should I find it then?
>>>
>>>So what I posted before is perhaps wrong. CONFIG_FAIRSCHEM _was set_ in
>>>.config file. But I got the error "fairsched_chwt: No such file or
>>>directory".
>>>
```

```
>>>Christian Kaiser
```

```
>>>
>>>
>>
>>--- ./kernel/fairsched.c.ve2350 2007-05-24 11:48:18.000000000 +0400
>>+++ ./kernel/fairsched.c 2007-05-25 16:36:36.000000000 +0400
>>@@ -618,6 +618,7 @@ static struct fairsched_node *fairsched_
>> struct fairsched_node *p;
>>
>> for_each_fairsched_node(p) {
>>+ printk(" fairsched_find: search %u, node: %d/%u\n", id, p->id,
>
> p->id);
>
>> if (p->id == id)
>> return p;
>> }
>>@@ -734,9 +735,13 @@ int do_fairsched_chwt(unsigned int id, u
>> if (weight < 1 || weight > FSCHWEIGHT_MAX)
>> return -EINVAL;
>>
>>+ printk("do_fairsched_chwt: node %u, w %u\n", id, weight);
>> node = fairsched_find(id);
>>- if (node == NULL)
>>+ if (node == NULL) {
>>+ printk("do_fairsched_chwt: ENOENT\n");
>> return -ENOENT;
>>+ }
>>+ printk("do_fairsched_chwt: found\n");
>>
>> spin_lock_irq(&fairsched_lock);
>> node->weight = weight;
>>diff --git a/src/lib/cpu.c b/src/lib/cpu.c
>>index 947d973..b0b53a7 100644
>>--- a/src/lib/cpu.c
>>+++ b/src/lib/cpu.c
>>@@ -74,7 +74,7 @@ static int set_cpulimit(envid_t veid, un
>>
>> static int set_cpuweight(envid_t veid, unsigned int cpuweight)
>> {
```

```
>>-
>>+ logger(0, 0, "vzctl: fairsched_chwt(%u, %u)", veid, cpuweight);
>> if (fairsched_chwt(veid, cpuweight)) {
>>     logger(-1, errno, "fairsched_chwt");
>>     return VZ_SETFSHD_ERROR;
>
>
```

```
> Christian Kaiser
> --
> IBM Deutschland Entwicklung GmbH
> Open Systems Firmware Development
> mail: ckaiser2@de.ibm.com
>
> IBM Deutschland Entwicklung GmbH
> Vorsitzender des Aufsichtsrats: Martin Jetter
```

```
> Registergericht: Amtsgericht Stuttgart, HRB 243294
>
>
>
```

```
--- ./include/asm-powerpc/systbl.h.ve2351 2007-05-24 11:48:17.000000000 +0400
+++ ./include/asm-powerpc/systbl.h 2007-05-25 17:58:34.000000000 +0400
@@ -304,7 +304,7 @@ SYSCALL_SPU(fchmodat)
SYSCALL_SPU(faccessat)
COMPAT_SYS_SPU(get_robust_list)
COMPAT_SYS_SPU(set_robust_list) /* 300 */
-SYS_SKIP(301, 399)
+SYS_SKIP(301, 400)
SYSCALL(ni_syscall)
SYS_SKIP_END()
SYSCALL(fairsched_mknod) /* 400 */
@@ -313,7 +313,7 @@ SYSCALL(fairsched_chwt)
SYSCALL(fairsched_mvpr)
SYSCALL(fairsched_rate)
SYSCALL(fairsched_vcpus)
-SYS_SKIP(406, 409)
+SYS_SKIP(406, 410)
SYSCALL(ni_syscall)
SYS_SKIP_END()
SYSCALL(getluid) /* 410 */
```
