

Re: [PATCH 0/2] NR\_CPUS: increase maximum NR\_CPUS to 4096

## Re: [PATCH 0/2] NR\_CPUS: increase maximum NR\_CPUS to 4096

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-04/msg03027.html>

---

- *From:* "Yinghai Lu" <yhlu.kernel@xxxxxxxx>
  - *Date:* Tue, 8 Apr 2008 14:53:51 -0700
- 

On Fri, Apr 4, 2008 at 6:30 PM, Mike Travis <travis@xxxxxxx> wrote:

\* Increases the limit of NR\_CPUS to 4096 and introduces a boolean called "MAXSMP" which when set (e.g. "allyesconfig") will set NR\_CPUS = 4096 and NODES\_SHIFT = 9 (512).

I've been running this config (4k NR\_CPUS, 512 Max Nodes) on an AMD box with 2 dual-cores and 4gb memory as well as an Intel box with 4 single-core cpus and 8Mb. I've also successfully booted it in a simulated 2cpus/1Gb environment.

Based on:

git://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux-2.6.git  
+ x86/latest .../x86/linux-2.6-x86.git  
+ sched-devel/latest .../mingo/linux-2.6-sched-devel.git

Signed-off-by: Mike Travis <travis@xxxxxxx>

got

-----[ cut here ]-----

WARNING: at kernel/sched\_fair.c:815 hrtick\_start\_fair+0x69/0x156()

Modules linked in:

Pid: 1, comm: swapper Not tainted

2.6.25-rc8-x86-latest.git-smp-01033-ga39ae31-dirty #77

Call Trace:

```
[<ffffffff802596ce>] warn_on_slowpath+0x67/0x8e
[<ffffffff8024b266>] hrtick_start_fair+0x69/0x156
[<ffffffff8024a619>] ? dequeue_entity+0x2a/0xf8
[<ffffffff8025547d>] dequeue_task_fair+0x5f/0x7e
[<ffffffff80248ea3>] dequeue_task+0x22/0x44
[<ffffffff80248efe>] deactivate_task+0x39/0x69
[<ffffffff80a57c31>] schedule+0x1b9/0x5c5
[<ffffffff80270dff>] ? autoremove_wake_function+0x20/0x5e
[<ffffffff80a58338>] schedule_timeout+0x31/0xd7
```

Re: [PATCH 0/2] NR\_CPUS: increase maximum NR\_CPUS to 4096

Re: [PATCH 0/2] NR\_CPUS: increase maximum NR\_CPUS to 4096

```
[<ffffffff8024db1a>] ? __wake_up+0x52/0x75
[<ffffffff80a578a3>] wait_for_common+0x103/0x189
[<ffffffff8024f29b>] ? default_wake_function+0x0/0x36
[<ffffffff80a57a62>] wait_for_completion+0x2b/0x41
[<ffffffff8026c0d2>] call_usermodehelper_exec+0x87/0xe5
[<ffffffff80561073>] kobject_uevent_env+0x3d0/0x424
[<ffffffff805610e5>] kobject_uevent+0x1e/0x34
[<ffffffff805f2dbd>] device_add+0x2f9/0x494
[<ffffffff805f2f80>] device_register+0x28/0x43
[<ffffffff80577abf>] pcie_port_device_register+0x3f1/0x43e
[<ffffffff80972165>] ? pcibios_set_master+0x8d/0xa8
[<ffffffff80a24a1c>] pcie_portdrv_probe+0x79/0xbb
[<ffffffff80574173>] pci_call_probe+0xe5/0x146
[<ffffffff80574331>] pci_device_probe+0x64/0xa2
[<ffffffff805f5946>] driver_probe_device+0xcf/0x16d
[<ffffffff8031d708>] ? sysfs_addrm_finish+0x2f/0x22b
[<ffffffff805f5b0b>] ? __driver_attach+0x0/0xbe
[<ffffffff805f5b79>] __driver_attach+0x6e/0xbe
[<ffffffff805f4999>] bus_for_each_dev+0x5e/0xa2
[<ffffffff805f571f>] driver_attach+0x2f/0x45
[<ffffffff805f5430>] bus_add_driver+0xc6/0x226
[<ffffffff805f45a4>] ? bus_put+0x29/0x3f
[<ffffffff805f5e66>] driver_register+0x6d/0xfc
[<ffffffff805745fa>] __pci_register_driver+0x62/0xb0
[<ffffffff818b25db>] pcie_portdrv_init+0x4a/0x72
[<ffffffff81890bcf>] kernel_init+0x1b4/0x340
[<ffffffff80225308>] child_rip+0xa/0x12
[<ffffffff81890a1b>] ? kernel_init+0x0/0x340
[<ffffffff802252fe>] ? child_rip+0x0/0x12
```

---[ end trace e26645195698f5cf ]---

BUG: unable to handle kernel NULL pointer dereference at 0000000000000148

IP: [<ffffffff8024b3cf>] pick\_next\_task\_fair+0x7c/0xbb

PGD 0

Oops: 0000 [1] SMP

CPU 28

Modules linked in:

Pid: 1, comm: swapper Not tainted

2.6.25-rc8-x86-latest.git-smp-01033-ga39ae31-dirty #77

RIP: 0010:[<ffffffff8024b3cf>] [<ffffffff8024b3cf>]

pick\_next\_task\_fair+0x7c/0xbb

RSP: 0018:fff81081cc5cd70 EFLAGS: 00010046

RAX: 0000000000000000 RBX: fff81383c21a280 RCX: 0000000000000000

RDX: fff81383c224080 RSI: fff81383c224080 RDI: 0000000063e15417

RBP: fff81081cc5cda0 R08: 0000000000000000 R09: fff81383c224108

R10: 0000000000000001 R11: 0000000000000000 R12: 0000000000000000

R13: fff81383c224080 R14: fff81383c224080 R15: 000000000000001c

FS: 0000000000000000(0000) GS: fff81401cc3c600(0000) knlGS:0000000000000000

CS: 0010 DS: 0018 ES: 0018 CR0: 000000008005003b

CR2: 0000000000000148 CR3: 000000000201000 CR4: 00000000000006e0

DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000

Re: [PATCH 0/2] NR\_CPUS: increase maximum NR\_CPUS to 4096

Re: [PATCH 0/2] NR\_CPUS: increase maximum NR\_CPUS to 4096

DR3: 0000000000000000 DR6: 00000000ffff0ff0 DR7: 0000000000000400  
Process swapper (pid: 1, threadinfo ffff81081cc5c000, task ffff81401cc52000)  
Stack: ffff81081cc5cda0 0000000063e15417 ffffffff80a81840 0000000000000000  
00000000fffeecfd ffff81383c224080 ffff81081cc5ce70 ffffffff80a57d28  
ffff81081cc5ce00 ffff81081cc5ce20 ffffffff81963080 ffffffff81963080

Call Trace:

[<ffffffffff80a57d28>] schedule+0x2b0/0x5c5  
[<ffffffffff80270dff>] ? autoremove\_wake\_function+0x20/0x5e  
[<ffffffffff80a58338>] schedule\_timeout+0x31/0xd7  
[<ffffffffff8024db1a>] ? \_\_wake\_up+0x52/0x75  
[<ffffffffff80a578a3>] wait\_for\_common+0x103/0x189  
[<ffffffffff8024f29b>] ? default\_wake\_function+0x0/0x36  
[<ffffffffff80a57a62>] wait\_for\_completion+0x2b/0x41  
[<ffffffffff8026c0d2>] call\_usermodehelper\_exec+0x87/0xe5  
[<ffffffffff80561073>] kobject\_uevent\_env+0x3d0/0x424  
[<ffffffffff805610e5>] kobject\_uevent+0x1e/0x34  
[<ffffffffff805f2dbd>] device\_add+0x2f9/0x494  
[<ffffffffff805f2f80>] device\_register+0x28/0x43  
[<ffffffffff80577abf>] pcie\_port\_device\_register+0x3f1/0x43e  
[<ffffffffff80972165>] ? pcibios\_set\_master+0x8d/0xa8  
[<ffffffffff80a24a1c>] pcie\_portdrv\_probe+0x79/0xbb  
[<ffffffffff80574173>] pci\_call\_probe+0xe5/0x146  
[<ffffffffff80574331>] pci\_device\_probe+0x64/0xa2  
[<ffffffffff805f5946>] driver\_probe\_device+0xcf/0x16d  
[<ffffffffff8031d708>] ? sysfs\_addrm\_finish+0x2f/0x22b  
[<ffffffffff805f5b0b>] ? \_\_driver\_attach+0x0/0xbe  
[<ffffffffff805f5b79>] \_\_driver\_attach+0x6e/0xbe  
[<ffffffffff805f4999>] bus\_for\_each\_dev+0x5e/0xa2  
[<ffffffffff805f571f>] driver\_attach+0x2f/0x45  
[<ffffffffff805f5430>] bus\_add\_driver+0xc6/0x226  
[<ffffffffff805f45a4>] ? bus\_put+0x29/0x3f  
[<ffffffffff805f5e66>] driver\_register+0x6d/0xfc  
[<ffffffffff805745fa>] \_\_pci\_register\_driver+0x62/0xb0  
[<ffffffffff818b25db>] pcie\_portdrv\_init+0x4a/0x72  
[<ffffffffff81890bcf>] kernel\_init+0x1b4/0x340  
[<ffffffffff80225308>] child\_rip+0xa/0x12  
[<ffffffffff81890a1b>] ? kernel\_init+0x0/0x340  
[<ffffffffff802252fe>] ? child\_rip+0x0/0x12

Code: 24 40 78 1c 8b 3d 36 05 b3 00 48 89 da be 00 04 00 00 e8 7a eb  
ff ff 49 39 c6 7f 04 4c 8b 63 48 4c 89 e6 4  
8 89 df e8 29 f1 ff ff <49> 8b 9c 24 48 01 00 00 48 85 db 75 a5 49 8d  
5c 24 c8 4c 89 ef

RIP [<ffffffffff8024b3cf>] pick\_next\_task\_fair+0x7c/0xbb

RSP [<ffff81081cc5cd70>]

CR2: 0000000000000148

---[ end trace e26645195698f5cf ]---

--

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in  
the body of a message to majordomo@xxxxxxxxxxxxxxxxxxx

Re: [PATCH 0/2] NR\_CPUS: increase maximum NR\_CPUS to 4096

Re: [PATCH 0/2] NR\_CPUS: increase maximum NR\_CPUS to 4096

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>