

---

Subject: [PATCH] pidns: make pid->level and pid\_ns->level unsigned

Posted by [Pavel Emelianov](#) on Tue, 26 Feb 2008 16:11:30 GMT

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

---

These values represent the nesting level of a namespace and pids living in it, and it's always non-negative.

Turning this from int to unsigned int saves some space in pid.c (11 bytes on x86 and 64 on ia64) by letting the compiler optimize the pid\_nr\_ns a bit. E.g. on ia64 this removes the sign extension calls, which compiler adds to optimize access to pid->numers[ns->level].

Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

---

```
diff --git a/include/linux/pid.h b/include/linux/pid.h
```

```
index c798081..03573e3 100644
```

```
--- a/include/linux/pid.h
```

```
+++ b/include/linux/pid.h
```

```
@ @ -60,7 +60,7 @ @ struct pid
```

```
/* lists of tasks that use this pid */
```

```
struct hlist_head tasks[PIDTYPE_MAX];
```

```
struct rcu_head rcu;
```

```
- int level;
```

```
+ unsigned int level;
```

```
struct upid numbers[1];
```

```
};
```

```
diff --git a/include/linux/pid_namespace.h b/include/linux/pid_namespace.h
```

```
index fcd61fa..caff528 100644
```

```
--- a/include/linux/pid_namespace.h
```

```
+++ b/include/linux/pid_namespace.h
```

```
@ @ -20,7 +20,7 @ @ struct pid_namespace {
```

```
int last_pid;
```

```
struct task_struct *child_reaper;
```

```
struct kmem_cache *pid_cache;
```

```
- int level;
```

```
+ unsigned int level;
```

```
struct pid_namespace *parent;
```

```
#ifdef CONFIG_PROC_FS
```

```
struct vfsmount *proc_mnt;
```

```
diff --git a/kernel/pid_namespace.c b/kernel/pid_namespace.c
```

```
index 6d792b6..cb17497 100644
```

```
--- a/kernel/pid_namespace.c
```

```
+++ b/kernel/pid_namespace.c
```

```
@ @ -66,7 +66,7 @ @ err_alloc:
```

```
    return NULL;
}

-static struct pid_namespace *create_pid_namespace(int level)
+static struct pid_namespace *create_pid_namespace(unsigned int level)
{
    struct pid_namespace *ns;
    int i;
```

---