

---

Subject: Re: [PATCH 8/15] Helpers to find the task by its numerical ids

Posted by Pavel Emelianov on Mon, 30 Jul 2007 06:15:22 GMT

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

---

Oleg Nesterov wrote:

> On 07/26, Pavel Emelyanov wrote:

```
>> +#define find_pid(pid) find_pid_ns(pid, &init_pid_ns)
>
> Again, I think find_pid() should use current's active ns, not
> init_pid_ns. Just grep for find_pid/find_task_by_pid.
>
>> --- linux-2.6.23-rc1-mm1.orig/kernel/pid.c 2007-07-26 16:34:45.000000000 +0400
>> +++ linux-2.6.23-rc1-mm1-7/kernel/pid.c 2007-07-26 16:36:37.000000000 +0400
>> @@ -204,19 +221,20 @@ static void delayed_put_pid(struct rcu_h
>>     goto out;
>> }
>>
>> -struct pid * fastcall find_pid(int nr)
>> +struct pid * fastcall find_pid_ns(int nr, struct pid_namespace *ns)
>> {
>>     struct hlist_node *elem;
>> - struct pid *pid;
>> + struct upid *pnr;
>> +
>> + hlist_for_each_entry_rcu(pnr, elem,
>> +     &pid_hash[pid_hashfn(nr, ns)], pid_chain)
>> + if (pnr->nr == nr && pnr->ns == ns)
>           ^^^^^^^^^^
> Aha, that is why we need upid->ns.
```

That's it :)

> I am a bit surprised we don't move the global pid\_hash into the  
> "struct pid\_namespace", this could speedup the search, and we  
> don't need upid->ns.

Hm... Worth thinking about, but this hash itself is large enough and  
its size depends on the node's number of pages, so we'll have

1. either to make per-namespace hash (much) smaller;
2. or to give (too) many memory for it.

```
>> -struct pid *find_ge_pid(int nr)
>> +struct pid *find_ge_pid(int nr, struct pid_namespace *ns)
>> {
>>     struct pid *pid;
>>
>>     do {
```

```
>> - pid = find_pid(nr);
>> + pid = find_pid_ns(nr, ns);
>> if (pid)
>> break;
>> - nr = next_pidmap(task_active_pid_ns(current), nr);
>> + nr = next_pidmap(ns, nr);
>> } while (nr > 0);
>>
>> return pid;
>
> This means we should fix the caller, next_tgid(), but this is done
> in 15/15.
```

Sorry :)

> Oleg.  
>  
>

Thank,  
Pavel

---