
Subject: [PATCH 7/16] Helpers to find the task by its numerical ids

Posted by [Pavel Emelianov](#) on Fri, 06 Jul 2007 08:07:02 GMT

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

When searching the task by numerical id one may need to find it using global pid (as it is done now in kernel) or by its virtual id, e.g. when sending a signal to a task from one namespace the sender will specify the task's virtual id.

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

```
fs/proc/base.c      | 2 +-
include/linux/pid.h  | 13 ++++++++
include/linux/sched.h | 31 ++++++++
kernel/pid.c         | 32 ++++++++
4 files changed, 58 insertions(+), 20 deletions(-)
```

--- ./fs/proc/base.c.ve6 2007-07-06 10:58:56.000000000 +0400

+++ ./fs/proc/base.c 2007-07-06 11:03:41.000000000 +0400

```
@ @ -2230,7 +2230,7 @ @ static struct task_struct *next_tgid(uns
    rcu_read_lock();
```

```
retry:
```

```
    task = NULL;
```

```
- pid = find_ge_pid(tgid);
```

```
+ pid = find_ge_pid(tgid, &init_pid_ns);
```

```
    if (pid) {
```

```
        tgid = pid->nr + 1;
```

```
        task = pid_task(pid, PIDTYPE_PID);
```

--- ./include/linux/pid.h.ve6 2007-07-06 11:03:27.000000000 +0400

+++ ./include/linux/pid.h 2007-07-06 11:03:27.000000000 +0400

```
@ @ -98,14 +98,23 @ @ extern struct pid_namespace init_pid_ns;
```

```
/*
```

```
 * look up a PID in the hash table. Must be called with the tasklist_lock
```

```
 * or rcu_read_lock() held.
```

```
+ *
```

```
+ * find_pid_ns() finds the pid in the namespace specified
```

```
+ * find_pid() find the pid by its global id, i.e. in the init namespace
```

```
+ * find_vpid() find the pid by its virtual id, i.e. in the current namespace
```

```
+ *
```

```
+ * see also find_task_by_pid() set in include/linux/sched.h
```

```
*/
```

```
-extern struct pid *FASTCALL(find_pid(int nr));
```

```
+extern struct pid *FASTCALL(find_pid_ns(int nr, struct pid_namespace *ns));
```

```
+
```

```
+#define find_vpid(pid) find_pid_ns(pid, current->nsproxy->pid_ns)
```

```
+#define find_pid(pid) find_pid_ns(pid, &init_pid_ns)
```

```

/*
 * Lookup a PID in the hash table, and return with it's count elevated.
 */
extern struct pid *find_get_pid(int nr);
-extern struct pid *find_ge_pid(int nr);
+extern struct pid *find_ge_pid(int nr, struct pid_namespace *);

extern struct pid *alloc_pid(void);
extern void FASTCALL(free_pid(struct pid *pid));
--- ./include/linux/sched.h.ve6 2007-07-06 11:03:27.000000000 +0400
+++ ./include/linux/sched.h 2007-07-06 11:03:27.000000000 +0400
@@ -1475,8 +1475,35 @@ extern struct task_struct init_task;

extern struct mm_struct init_mm;

-#define find_task_by_pid(nr) find_task_by_pid_type(PIDTYPE_PID, nr)
-extern struct task_struct *find_task_by_pid_type(int type, int pid);
+extern struct pid_namespace init_pid_ns;
+
+/*
+ * find a task by one of its numerical ids
+ *
+ * find_task_by_pid_type_ns():
+ *   it is the most generic call - it finds a task by all id,
+ *   type and namespace specified
+ * find_task_by_pid_ns():
+ *   finds a task by its pid in the specified namespace
+ * find_task_by_pid_type():
+ *   finds a task by its global id with the specified type, e.g.
+ *   by global session id
+ * find_task_by_pid():
+ *   finds a task by its global pid
+ *
+ * see also find_pid() etc in include/linux/pid.h
+ */
+
+extern struct task_struct *find_task_by_pid_type_ns(int type, int pid,
+ struct pid_namespace *ns);
+
+#define find_task_by_pid_ns(nr, ns) \
+ find_task_by_pid_type_ns(PIDTYPE_PID, nr, ns)
+#define find_task_by_pid_type(type, nr) \
+ find_task_by_pid_type_ns(type, nr, &init_pid_ns)
+#define find_task_by_pid(nr) \
+ find_task_by_pid_type(PIDTYPE_PID, nr)
+
extern void __set_special_pids(pid_t session, pid_t pgrp);

```

```

/* per-UID process charging. */
--- ./kernel/pid.c.ve6 2007-07-06 11:03:27.000000000 +0400
+++ ./kernel/pid.c 2007-07-06 11:03:27.000000000 +0400
@@ -238,19 +238,20 @@ out_free:
    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 pid_number *pnr;
+
+ hlist_for_each_entry_rcu(pnr, elem,
+ &pid_hash[pid_hashfn(nr)], pid_chain)
+ if (pnr->nr == nr && pnr->ns == ns)
+ return container_of(pnr, struct pid,
+ numbers[ns->level]);

- hlist_for_each_entry_rcu(pid, elem,
- &pid_hash[pid_hashfn(nr)], pid_chain) {
- if (pid->nr == nr)
- return pid;
- }
    return NULL;
}
-EXPORT_SYMBOL_GPL(find_pid);
+EXPORT_SYMBOL_GPL(find_pid_ns);

/*
 * attach_pid() must be called with the tasklist_lock write-held.
@@ -310,12 +311,13 @@ struct task_struct * fastcall pid_task(s
/*
 * Must be called under rcu_read_lock() or with tasklist_lock read-held.
 */
-struct task_struct *find_task_by_pid_type(int type, int nr)
+struct task_struct *find_task_by_pid_type_ns(int type, int nr,
+ struct pid_namespace *ns)
{
- return pid_task(find_pid(nr), type);
+ return pid_task(find_pid_ns(nr, ns), type);
}

-EXPORT_SYMBOL(find_task_by_pid_type);
+EXPORT_SYMBOL(find_task_by_pid_type_ns);

```

```

struct pid *get_task_pid(struct task_struct *task, enum pid_type type)
{
@@ -342,7 +344,7 @@ struct pid *find_get_pid(pid_t nr)
    struct pid *pid;

    rcu_read_lock();
- pid = get_pid(find_pid(nr));
+ pid = get_pid(find_vpid(nr));
    rcu_read_unlock();

    return pid;
@@ -361,15 +363,15 @@ pid_t pid_nr_ns(struct pid *pid, struct
    *
    * If there is a pid at nr this function is exactly the same as find_pid.
    */
-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(current->nsproxy->pid_ns, nr);
+ nr = next_pidmap(ns, nr);
    } while (nr > 0);

    return pid;

```

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