
Subject: [PATCH 11/28] [PREP 11/14] Add support for multiple hash tables in pid.c

Posted by Pavel Emelianov on Fri, 15 Jun 2007 16:09:35 GMT

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

It turned out that virtual pids require two hash tables for pid searching.
E.g. flat model hashed pid by its global and virtual ids, multilevel model
uses one hash to find the pid by number, and the other one to find the
number by the struct pid.

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

pid.c | 18 ++++++-----

1 files changed, 17 insertions(+), 1 deletion(-)

```
--- ./kernel/pid.c.pidmorecaches 2007-06-15 15:09:29.000000000 +0400
+++ ./kernel/pid.c 2007-06-15 15:11:14.000000000 +0400
@@ -54,6 +54,17 @@ static inline int mk_pid(struct pid_name
#define find_next_offset(map, off) \
    find_next_zero_bit((map)->page, BITS_PER_PAGE, off)

+#ifdef CONFIG_PID_NS
+/*
+ * pid namespaces will require the additional hash to store the
+ * pid-to-namespace relations. so declare it here and define a hash
+ * fun of two arguments - nr and the namespce
+ */
+static struct hlist_head *pid_hash2;
+#define pid_ehashfn(nr, ns) hash_long((unsigned long)nr + (unsigned long)ns, \
+    pidhash_shift)
#endif
+
/*
 * PID-map pages start out as NULL, they get allocated upon
 * first use and are never deallocated. This way a low pid_max
@@ -420,12 +431,17 @@ void __init pidhash_init(void)
    printk("PID hash table entries: %d (order: %d, %Zd bytes)\n",
    pidhash_size, pidhash_shift,
    pidhash_size * sizeof(struct hlist_head));

-#ifdef CONFIG_PID_NS
+pidhash_size *= 2;
#endif
    pid_hash = alloc_bootmem(pidhash_size * sizeof(*pid_hash));
    if (!pid_hash)
        panic("Could not alloc pidhash!\n");
    for (i = 0; i < pidhash_size; i++)
```

```
INIT_HLIST_HEAD(&pid_hash[i]);  
+#ifdef CONFIG_PID_NS  
+ pid_hash2 = pid_hash + (pidhash_size / 2);  
+#endif  
}
```

```
void __init pidmap_init(void)
```

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
