Subject: [PATCH 2.6.24-rc8-mm1 04/15] IPC/semaphores: move the rwmutex handling inside semctl_down
Posted by Pierre Peiffer on Tue, 29 Jan 2008 16:02:33 GMT
View Forum Message <> Reply to Message

From: Pierre Peiffer <pierre.peiffer@bull.net>

semctl_down is called with the rwmutex (the one which protects the
list of ipcs) taken in write mode.
This patch moves this rwmutex taken in write-mode inside semctl_down.
This has the advantages of reducing a little bit the window during
which this rwmutex is taken, clarifying sys_semctl, and finally of
having a coherent behaviour with [shm|msg]ctl_down

Signed-off-by: Pierre Peiffer <pierre.peiffer@bull.net>
Acked-by: Serge Hallyn <serue@us.ibm.com>
---

 ipc/sem.c | 24 ++++++++++++++++----------
 1 file changed, 13 insertions(+), 11 deletions(-)

Index: b/ipc/sem.c
===
--- a/ipc/sem.c
+++ b/ipc/sem.c
@@ -877,6 +877,11 @@ static inline unsigned long copy_semid_f
  }
 }

+/*
+ * This function handles some semctl commands which require the rw_mutex
+ * to be held in write mode.
+ * NOTE: no locks must be held, the rw_mutex is taken inside this function.
+ */
 static int semctl_down(struct ipc_namespace *ns, int semid, int semnum,
    int cmd, int version, union semun arg)
 {
@@ -889,9 +894,12 @@ static int semctl_down(struct ipc_namesp
    if(copy_semid_from_user (&setbuf, arg.buf, version))
     return -EFAULT;
   }
+ down_write(&sem_ids(ns).rw_mutex);
   sma = sem_lock_check_down(ns, semid);
- if (IS_ERR(sma))
- return PTR_ERR(sma);
+ if (IS_ERR(sma)) {
+ err = PTR_ERR(sma);
+ goto out_up;

```
+ }

  ipcp = &sma->sem_perm;

@@ -917,26 +925,22 @@ static int semctl_down(struct ipc_namesp
  switch(cmd){
  case IPC_RMID:
   freeary(ns, ipcp);
-  err = 0;
-  break;
+  goto out_up;
  case IPC_SET:
   ipcp->uid = setbuf.uid;
   ipcp->gid = setbuf.gid;
   ipcp->mode = (ipcp->mode & ~S_IRWXUGO)
     | (setbuf.mode & S_IRWXUGO);
   sma->sem_ctime = get_seconds();
-  sem_unlock(sma);
-  err = 0;
   break;
  default:
-  sem_unlock(sma);
   err = -EINVAL;
-  break;
  }
- return err;

 out_unlock:
  sem_unlock(sma);
+out_up:
+ up_write(&sem_ids(ns).rw_mutex);
  return err;
 }

@@ -970,9 +974,7 @@ asmlinkage long sys_semctl (int semid, i
   return err;
  case IPC_RMID:
  case IPC_SET:
-  down_write(&sem_ids(ns).rw_mutex);
   err = semctl_down(ns,semid,semnum,cmd,version,arg);
-  up_write(&sem_ids(ns).rw_mutex);
   return err;
  default:
   return -EINVAL;

--
Pierre Peiffer
```

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