
Subject: [PATCH v7 00/10] IPC: checkpoint/restore in userspace enhancements
Posted by [Stanislav Kinsbursky](#) on Thu, 18 Oct 2012 10:22:21 GMT
[View Forum Message](#) <> [Reply to Message](#)

v7:

- 1) Added comments in places, when behaviour is not obvious.
- 2) Compilation fixed for compat layer
- 3) test updated to use existent header files (instead of hard-coding new defines in case of absence).
- 4) comment fixed in qlge driver

v6:

- 1) rebased on 3.7-rc1

v5:

- 1) Several define-dependent compile bugs fixed
- 2) IPC message copy test updated
- 3) A couple of minor fixes.
- 4) Qlogic driver update: rename of its internal SEM_SET define into SEM_INIT (compile error).

v4:

- 1) If MSG_COPY flag is specified, then "mtype" is not a type, but message number to copy.
- 2) MSG_SET_COPY logic for sys_msgctl() was removed.

v3:

- 1) Copy messages to user-space under spinlock was replaced by allocation of dummy message before queue lock and then copy of desired message to the dummy one instead of unlinking it from queue list.
I.e. the message queue copy logic was changed: messages can be retrived one by one (instead of receiving of the whole list at once).

This patch set is aimed to provide additional functionality for all IPC objects,
which is required for migration of these objects by user-space checkpoint/restore utils (CRIU).

The main problem here was impossibility to set up object id. This patch set solves the problem in two steps:

- 1) Makes it possible to create new object (shared memory, semaphores set or messages queue) with ID, equal to passed key.
- 2) Makes it possible to change existent object key.

Another problem was to peek messages from queues without deleting them. This was achived by introducing of new MSG_COPY flag for sys_msgrcv(). If MSG_COPY flag is set, then msgtyp is interpreted as message number.

The following series implements...

Stanislav Kinsbursky (10):

- ipc: remove forced assignment of selected message
- ipc: "use key as id" functionality for resource get system call introduced
- ipc: segment key change helper introduced
- ipc: add new SHM_SET command for sys_shmctl() call
- ipc: add new MSG_SET command for sys_msgctl() call
- qlge driver: rename internal SEM_SET macro to SEM_INIT
- ipc: add new SEM_SET command for sys_semctl() call
- IPC: message queue receive cleanup
- IPC: message queue copy feature introduced
- test: IPC message queue copy feture test

```
drivers/net/ethernet/qlogic/qlge/qlge.h      |  4
drivers/net/ethernet/qlogic/qlge/qlge_main.c | 16 +-
include/linux/msg.h                          |  5 -
include/uapi/linux/ipc.h                    |  1
include/uapi/linux/msg.h                    |  2
include/uapi/linux/sem.h                    |  1
include/uapi/linux/shm.h                    |  1
ipc/compat.c                                | 54 +++---
ipc/msg.c                                    | 117 ++++++++----
ipc/msgutil.c                               | 38 ++++
ipc/sem.c                                    | 15 +-
ipc/shm.c                                    | 18 ++
ipc/util.c                                  | 67 ++++++--
ipc/util.h                                  |  6 +
security/selinux/hooks.c                    |  3
security/smack/smack_lsm.c                  |  3
tools/testing/selftests/ipc/Makefile         | 25 +++
tools/testing/selftests/ipc/msgqueue.c      | 231 ++++++++
18 files changed, 526 insertions(+), 81 deletions(-)
create mode 100644 tools/testing/selftests/ipc/Makefile
create mode 100644 tools/testing/selftests/ipc/msgqueue.c
```
