
Subject: [PATCH 2/4] userns: let copy_ipcs handle setting ipc_ns->user_ns
Posted by [serge](#) on Mon, 21 Feb 2011 04:02:30 GMT

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

To do that, we have to pass in the task_struct of the task which will own the ipc_ns, so we can assign its user_ns.

Signed-off-by: Serge E. Hallyn <serge.hallyn@canonical.com>

```
---  
include/linux/ipc_namespace.h |  8 ++++++--  
ipc/namespace.c           | 12 ++++++----  
kernel/nsproxy.c          |  7 +-----  
3 files changed, 13 insertions(+), 14 deletions(-)  
  
diff --git a/include/linux/ipc_namespace.h b/include/linux/ipc_namespace.h  
index 46d2eb4..9974429 100644  
--- a/include/linux/ipc_namespace.h  
+++ b/include/linux/ipc_namespace.h  
@@ -92,7 +92,8 @@ static inline int mq_init_ns(struct ipc_namespace *ns) { return 0; }  
#endif  
  
#if defined(CONFIG_IPC_NS)  
-extern struct ipc_namespace *copy_ipcs(unsigned long flags,  
+extern struct ipc_namespace *copy_ipcs(struct task_struct *tsk,  
+    unsigned long flags,  
    struct ipc_namespace *ns);  
static inline struct ipc_namespace *get_ipc_ns(struct ipc_namespace *ns)  
{  
@@ -103,8 +104,9 @@ static inline struct ipc_namespace *get_ipc_ns(struct ipc_namespace  
*ns)  
  
extern void put_ipc_ns(struct ipc_namespace *ns);  
#else  
-static inline struct ipc_namespace *copy_ipcs(unsigned long flags,  
- struct ipc_namespace *ns)  
+static inline struct ipc_namespace *copy_ipcs(struct task_struct *tsk,  
+    unsigned long flags,  
+    struct ipc_namespace *ns)  
{  
    if (flags & CLONE_NEWIPC)  
        return ERR_PTR(-EINVAL);  
diff --git a/ipc/namespace.c b/ipc/namespace.c  
index aa18899..ee84882 100644  
--- a/ipc/namespace.c  
+++ b/ipc/namespace.c  
@@ -15,7 +15,8 @@  
  
#include "util.h"
```

```

-static struct ipc_namespace *create_ipc_ns(struct ipc_namespace *old_ns)
+static struct ipc_namespace *create_ipc_ns(struct task_struct *tsk,
+    struct ipc_namespace *old_ns)
{
    struct ipc_namespace *ns;
    int err;
@@ -44,17 +45,18 @@ static struct ipc_namespace *create_ipc_ns(struct ipc_namespace
*old_ns)
    ipcns_notify(IPCNS_CREATED);
    register_ipcns_notifier(ns);

- ns->user_ns = old_ns->user_ns;
- get_user_ns(ns->user_ns);
+ ns->user_ns = get_user_ns(task_cred_xxx(tsk, user)->user_ns);

    return ns;
}

-struct ipc_namespace *copy_ipcs(unsigned long flags, struct ipc_namespace *ns)
+struct ipc_namespace *copy_ipcs(struct task_struct *tsk,
+    unsigned long flags,
+    struct ipc_namespace *ns)
{
    if (!(flags & CLONE_NEWIPC))
        return get_ipc_ns(ns);
- return create_ipc_ns(ns);
+ return create_ipc_ns(tsk, ns);
}

/*
diff --git a/kernel/nsproxy.c b/kernel/nsproxy.c
index ffa6b67..b905ecc 100644
--- a/kernel/nsproxy.c
+++ b/kernel/nsproxy.c
@@ -75,16 +75,11 @@ static struct nsproxy *create_new_namespaces(unsigned long flags,
    goto out_uts;
}

- new_nsp->ipc_ns = copy_ipcs(flags, tsk->nsproxy->ipc_ns);
+ new_nsp->ipc_ns = copy_ipcs(tsk, flags, tsk->nsproxy->ipc_ns);
    if (IS_ERR(new_nsp->ipc_ns)) {
        err = PTR_ERR(new_nsp->ipc_ns);
        goto out_ipc;
    }
- if (new_nsp->ipc_ns != tsk->nsproxy->ipc_ns) {
-     put_user_ns(new_nsp->ipc_ns->user_ns);
-     new_nsp->ipc_ns->user_ns = task_cred_xxx(tsk, user)->user_ns;

```

```
- get_user_ns(new_nsp->ipc_ns->user_ns);
- }

new_nsp->pid_ns = copy_pid_ns(flags, task_active_pid_ns(tsk));
if (IS_ERR(new_nsp->pid_ns)) {
--
```

1.7.0.4

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