
Subject: [PATCH][RFC] Cleanup in namespaces unsharing
Posted by [Pavel Emelianov](#) on Fri, 08 Jun 2007 09:09:16 GMT
[View Forum Message](#) <> [Reply to Message](#)

Currently we have two funtions to copy the namespaces:
copy_namespaces() and unshare_nsproxy_namespaces(). The
second one checks for unsupported functionality with

```
#ifndef CONFIG_IPC_NS
if (unshare_flags & CLONE_NEWIPC)
    return -EINVAL;
#endif
```

-like constructions, while the first one does not. One
of the side effects of this is that clone() with the
CLONE_NEWXXX set will return 0 if the kernel doesn't
support XXX namespaces thus confusing the user-level.

The proposal is to make these calls clean from the ifdefs
and move these checks into each namespaces' stubs. This
will make the code cleaner and (!) return -EINVAL from
fork() in case the desired namespaces are not supported.

Did I miss something in the design or this patch worth merging?

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

```
diff --git a/include/linux/ipc.h b/include/linux/ipc.h
index 7c8c6d8..b5aed71 100644
--- a/include/linux/ipc.h
+++ b/include/linux/ipc.h
@@ -100,6 +100,9 @@ extern struct ipc_namespace *copy_ipcs(u
 static inline struct ipc_namespace *copy_ipcs(unsigned long flags,
       struct ipc_namespace *ns)
{
+ if (flags & CLONE_NEWIPC)
+ ns = ERR_PTR(-EINVAL);
+
 return ns;
}
#endif

diff --git a/include/linux/utsname.h b/include/linux/utsname.h
index f8d3b32..230706e 100644
--- a/include/linux/utsname.h
+++ b/include/linux/utsname.h
@@ -60,6 +60,9 @@ static inline void put_uts_ns(struct uts
```

```

static inline struct uts_namespace *copy_utsname(int flags,
                                              struct uts_namespace *ns)
{
+ if (flags & CLONE_NEWUTS)
+   ns = ERR_PTR(-EINVAL);
+
  return ns;
}

diff --git a/kernel/nsproxy.c b/kernel/nsproxy.c
index 1bc4b55..ef26615 100644
--- a/kernel/nsproxy.c
+++ b/kernel/nsproxy.c
@@ -157,16 +157,6 @@ int unshare_nsproxy_namespaces(unsigned
  if (!(unshare_flags & (CLONE_NEWNS | CLONE_NEWUTS | CLONE_NEWIPC)))
    return 0;

-#ifndef CONFIG_IPC_NS
-if (unshare_flags & CLONE_NEWIPC)
-  return -EINVAL;
-#endif
-
-#ifndef CONFIG_UTS_NS
-if (unshare_flags & CLONE_NEWUTS)
-  return -EINVAL;
-#endif
-
  if (!capable(CAP_SYS_ADMIN))
    return -EPERM;

```
