
Subject: [patch 04/10] add "permit user mounts" flag to namespaces

Posted by Miklos Szeredi on Thu, 12 Apr 2007 16:45:45 GMT

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

From: Miklos Szeredi <mszeredi@suse.cz>

If MNT_NS_PERMIT_USERMOUNTS flag is not set for the current namespace, then unprivileged mounts will be denied.

By default this flag is cleared in all namespaces.

Signed-off-by: Miklos Szeredi <mszeredi@suse.cz>

Index: linux/fs/namespace.c

```
=====
--- linux.orig/fs/namespace.c 2007-04-12 16:50:16.000000000 +0200
+++ linux/fs/namespace.c 2007-04-12 16:50:17.000000000 +0200
@@ -1526,6 +1526,19 @@ dput_out:
     return retval;
 }

+static struct mnt_namespace *alloc_ns(void)
+{
+ struct mnt_namespace *ns;
+
+ ns = kzalloc(sizeof(struct mnt_namespace), GFP_KERNEL);
+ if (ns) {
+ atomic_set(&ns->count, 1);
+ INIT_LIST_HEAD(&ns->list);
+ init_waitqueue_head(&ns->poll);
+ }
+ return ns;
+}
+
+/*
+ * Allocate a new namespace structure and populate it with contents
+ * copied from the namespace of the passed in task structure.
@@ -1537,15 +1550,10 @@ static struct mnt_namespace *dup_mnt_ns(
     struct vfsmount *rootmnt = NULL, *pwdmnt = NULL, *altrootmnt = NULL;
     struct vfsmount *p, *q;

- new_ns = kmalloc(sizeof(struct mnt_namespace), GFP_KERNEL);
+ new_ns = alloc_ns();
     if (!new_ns)
         return NULL;

- atomic_set(&new_ns->count, 1);
```

```

- INIT_LIST_HEAD(&new_ns->list);
- init_waitqueue_head(&new_ns->poll);
- new_ns->event = 0;
-
  down_write(&namespace_sem);
  /* First pass: copy the tree topology */
  new_ns->root = copy_tree(mnt_ns->root, mnt_ns->root->mnt_root,
@@ -1860,13 +1868,10 @@ static void __init init_mount_tree(void)
  mnt = do_kern_mount("rootfs", 0, "rootfs", NULL);
  if (IS_ERR(mnt))
    panic("Can't create rootfs");
- ns = kmalloc(sizeof(*ns), GFP_KERNEL);
+ ns = alloc_ns();
  if (!ns)
    panic("Can't allocate initial namespace");
- atomic_set(&ns->count, 1);
- INIT_LIST_HEAD(&ns->list);
- init_waitqueue_head(&ns->poll);
- ns->event = 0;
+
  list_add(&mnt->mnt_list, &ns->list);
  ns->root = mnt;
  mnt->mnt_ns = ns;

```

Index: linux/include/linux/mnt_namespace.h

```

=====
--- linux.orig/include/linux/mnt_namespace.h 2007-04-12 16:50:02.000000000 +0200
+++ linux/include/linux/mnt_namespace.h 2007-04-12 16:50:17.000000000 +0200
@@ -6,12 +6,16 @@
#include <linux/sched.h>
#include <linux/nsproxy.h>

+/* mnt_namespace flags */
+#define MNT_NS_PERMIT_USERMOUNTS (1 << 0)
+
struct mnt_namespace {
  atomic_t count;
  struct vfsmount * root;
  struct list_head list;
  wait_queue_head_t poll;
  int event;
+ int flags;
};

extern struct mnt_namespace *copy_mnt_ns(int, struct mnt_namespace *,

--

```

Containers mailing list

