Subject: [patch 10/10] unprivileged mounts: add "no submounts" flag
Posted by Miklos Szeredi on Fri, 27 Apr 2007 12:04:27 GMT
View Forum Message <> Reply to Message

From: Miklos Szeredi <mszeredi@suse.cz>

Add a new mount flag "nomnt", which denies submounts for the owner.
This would be useful, if we want to support traditional /etc/fstab
based user mounts.

In this case mount(8) would still have to be suid-root, to check the
mountpoint against the user/users flag in /etc/fstab, but /etc/mtab
would no longer be mandatory for storing the actual owner of the
mount.

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

Index: linux/fs/namespace.c
===
--- linux.orig/fs/namespace.c 2007-04-27 12:57:11.000000000 +0200
+++ linux/fs/namespace.c 2007-04-27 12:57:14.000000000 +0200
@@ -449,6 +449,7 @@ static int show_vfsmnt(struct seq_file *
  { MNT_NOATIME, ",noatime" },
  { MNT_NODIRATIME, ",nodiratime" },
  { MNT_RELATIME, ",relatime" },
+ { MNT_NOMNT, ",nomnt" },
  { 0, NULL }
 };
 struct proc_fs_info *fs_infop;
@@ -806,6 +807,9 @@ static bool permit_mount(struct nameidat
 if (S_ISLNK(inode->i_mode))
  return false;

+ if (nd->mnt->mnt_flags & MNT_NOMNT)
+ return false;
+
 if (!is_mount_owner(nd->mnt, current->fsuid))
  return false;

@@ -1575,9 +1579,11 @@ long do_mount(char *dev_name, char *dir_
  mnt_flags |= MNT_NODIRATIME;
 if (flags & MS_RELATIME)
  mnt_flags |= MNT_RELATIME;
+ if (flags & MS_NOMNT)
+ mnt_flags |= MNT_NOMNT;

 flags &= ~(MS_NOSUID | MS_NOEXEC | MS_NODEV | MS_ACTIVE |

```
-     MS_NOATIME | MS_NODIRATIME | MS_RELATIME);
+     MS_NOATIME | MS_NODIRATIME | MS_RELATIME | MS_NOMNT);

   /* ... and get the mountpoint */
   retval = path_lookup(dir_name, LOOKUP_FOLLOW, &nd);
Index: linux/include/linux/fs.h
===================================================================
--- linux.orig/include/linux/fs.h 2007-04-27 12:57:11.000000000 +0200
+++ linux/include/linux/fs.h 2007-04-27 12:57:14.000000000 +0200
@@ -128,6 +128,7 @@ extern int dir_notify_enable;
 #define MS_SHARED (1<<20) /* change to shared */
 #define MS_RELATIME (1<<21) /* Update atime relative to mtime/ctime. */
 #define MS_SETUSER (1<<22) /* set mnt_uid to current user */
+#define MS_NOMNT (1<<23) /* don't allow unprivileged submounts */
 #define MS_ACTIVE (1<<30)
 #define MS_NOUSER (1<<31)

Index: linux/include/linux/mount.h
===================================================================
--- linux.orig/include/linux/mount.h 2007-04-27 12:57:01.000000000 +0200
+++ linux/include/linux/mount.h 2007-04-27 12:57:14.000000000 +0200
@@ -28,6 +28,7 @@ struct mnt_namespace;
 #define MNT_NOATIME 0x08
 #define MNT_NODIRATIME 0x10
 #define MNT_RELATIME 0x20
+#define MNT_NOMNT 0x40

 #define MNT_SHRINKABLE 0x100
 #define MNT_USER 0x200

--

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