
Subject: Re: [patch 9/9] unprivileged mounts: add "no submounts" flag
Posted by [serue](#) on Mon, 14 Jan 2008 23:39:10 GMT

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

Quoting Miklos Szeredi (miklos@szeredi.hu):

> 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.

Ah, I see, so the floppy drive could be mounted as a MNT_NOMNT but
MNT_USER mount with mnt_owner set. Makes sense. I'd ask for a better
name than 'nomnt', but I can't think of one myself.

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

Acked-by: Serge Hallyn <serue@us.ibm.com>

> ---
>
> Index: linux/fs/namespace.c
> ======
> --- linux.orig/fs/namespace.c 2008-01-04 13:49:52.000000000 +0100
> +++ linux/fs/namespace.c 2008-01-04 13:50:28.000000000 +0100
> @@ -694,6 +694,7 @@ 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;
> @@ -1044,6 +1045,9 @@ static bool permit_mount(struct nameidata
> if (S_ISLNK(inode->i_mode))
> return false;
>
> + if (nd->path.mnt->mnt_flags & MNT_NOMNT)
> + return false;
> +
> if (!is_mount_owner(nd->path.mnt, current->fsuid))
> return false;
>

```

> @@ -1888,9 +1892,11 @@ long do_mount(char *dev_name, char *dir_
>   mnt_flags |= MNT_RELATIME;
>   if (flags & MS_RDONLY)
>     mnt_flags |= MNT_READONLY;
> + if (flags & MS_NOMNT)
> +   mnt_flags |= MNT_NOMNT;
>
> - flags &= ~(MS_NOSUID | MS_NOEXEC | MS_NODEV | MS_ACTIVE |
> -   MS_NOATIME | MS_NODIRATIME | MS_RELATIME| MS_KERNMOUNT);
> + flags &= ~(MS_NOSUID | MS_NOEXEC | MS_NODEV | MS_ACTIVE | MS_NOATIME |
> +   MS_NODIRATIME | MS_RELATIME | MS_KERNMOUNT | 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 2008-01-04 13:49:12.000000000 +0100
> +++ linux/include/linux/fs.h 2008-01-04 13:49:58.000000000 +0100
> @@ -130,6 +130,7 @@ extern int dir_notify_enable;
> #define MS_KERNMOUNT (1<<22) /* this is a kern_mount call */
> #define MS_I_VERSION (1<<23) /* Update inode i_version field */
> #define MS_SETUSER (1<<24) /* set mnt_uid to current user */
> +#define MS_NOMNT (1<<25) /* 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 2008-01-04 13:45:45.000000000 +0100
> +++ linux/include/linux/mount.h 2008-01-04 13:49:58.000000000 +0100
> @@ -30,6 +30,7 @@ struct mnt_namespace;
> #define MNT_NODIRATIME 0x10
> #define MNT_RELATIME 0x20
> #define MNT_READONLY 0x40 /* does the user want this to be r/o? */
> +#define MNT_NOMNT 0x80
>
> #define MNT_SHRINKABLE 0x100
> #define MNT_IMBALANCED_WRITE_COUNT 0x200 /* just for debugging */
>
> --

```

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