

Re: [NFS] [PATCH 2/7] NFS: if ATTR\_KILL\_S\*ID bits are set, then skip mode change

## Re: [NFS] [PATCH 2/7] NFS: if ATTR\_KILL\_S\*ID bits are set, then skip mode change

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-09/msg03885.html>

---

- *From:* Greg Banks <gnb@xxxxxxx>
  - *Date:* Sat, 15 Sep 2007 00:40:33 +1000
- 

On Fri, Sep 14, 2007 at 09:38:46AM -0400, Jeff Layton wrote:

On Fri, 14 Sep 2007 23:09:24 +1000  
Greg Banks <gnb@xxxxxxx> wrote:

On Fri, Sep 14, 2007 at 07:02:58AM -0400, Jeff Layton wrote:

On Fri, 14 Sep 2007 20:25:45 +1000  
Greg Banks <gnb@xxxxxxx> wrote:

I'm curious about the reasons behind this change. You mention credential issues; how exactly is it that you have the correct creds to perform a WRITE rpc but not a SETATTR rpc?

Consider this case. user1 and user2 are both members of group "allusers":

```
user1$ echo foo > foo
user1$ chgrp allusers foo
user1$ chmod 04770 foo
user2$ echo bar >> foo
```

On most local filesystems, this would work correctly. The end result would be a file with mode 0770 and the expected contents. On NFS though, the write by user2 fails. When the write is attempted, the kernel tries to squash the setuid bit using the credentials of user2,

Re: [NFS] [PATCH 2/7] NFS: if ATTR\_KILL\_S\*ID bits are set, then skip mode change

who's not allowed to change the mode. The write then fails because the setattr fails.

Ok, I ran an experiment and I see this failure mode.

So the SETATTR rpc is really a side effect of the client kernel's behaviour and not an operation directly requested by the user process on the client. Is there any reason why that rpc needs to have user2's creds? Why not do the rpc with a fake set of creds with uid and gid set to the uid and gid of the file, in this case user1/allusers ? That way the rpc will most likely pass the server's permission check.

That might work in some cases, but there are many where it wouldn't...

Suppose user1 here is root and all of the user1 operations are being done on the server. If the server has root squashing enabled, then user2's operation would still fail.

In that case, user1's operations would also fail, which is even more serious a problem. Also arguably you actually \*want\* writes by a nonroot user to a setuid root executable to fail ;-)

Another problem:

Suppose we're using gssapi. There's no guarantee that the client will have the proper credentials to fake up a call as user1 (you might need user1 krb5 tickets, etc).

Yes, good point. You could use the root creds, except for root squashing. Ok, you convinced me.

Greg.

--

Greg Banks, R&D Software Engineer, SGI Australian Software Group.

Apparently, I'm Bedevere. Which MPHG character are you?

I don't speak for SGI.

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>

Re: [NFS] [PATCH 2/7] NFS: if ATTR\_KILL\_S\*ID bits are set, then skip mode change