Subject: Re: [PATCH] struct file leakage
Posted by Trond Myklebust on Tue, 11 Jul 2006 12:04:06 GMT
View Forum Message <> Reply to Message

On Mon, 2006-07-10 at 03:05 -0700, Andrew Morton wrote:
> On Mon, 10 Jul 2006 13:05:35 +0400
> Kirill Korotaev <dev@sw.ru> wrote:
>
> > Hello!
> >
> > Andrew, this is a patch from Alexey Kuznetsov for 2.6.16.
> > I believe 2.6.17 still has this leak.
> >
> > ---------- -
> >
> > 2.6.16 leaks like hell. While testing, I found massive leakage
> > (reproduced in openvz) in:
> >
> > *filp
> > *size-4096
> >
> > And 1 object leaks in
> > *size-32
> > *size-64
> > *size-128
> >
> >
> > It is the fix for the first one. filp leaks in the bowels
> > of namei.c.
> >
> > Seems, size-4096 is file table leaking in expand_fdtables.
>
> I suspect that's been there for a long time.
>
> > I have no idea what are the rest and why they show only
> > accompaniing another leaks. Some debugging structs?
>
> I don't understand this.  Are you implying that there are other bugs.
>
> > Signed-Off-By: Alexey Kuznetsov <kuznet@ms2.inr.ac.ru>
> > CC: Kirill Korotaev <dev@openvz.org>
> >
>
> > --- linux-2.6.16-w/fs/namei.c 2006-07-10 11:43:11.000000000 +0400
> > +++ linux-2.6.16/fs/namei.c 2006-07-10 11:53:36.000000000 +0400
> > @@ -1774,8 +1774,15 @@ do_link:
> >   if (error)
> >    goto exit_dput;

> >   error = __do_follow_link(&path, nd);
> > - if (error)
> > + if (error) {
> > +  /* Does someone understand code flow here? Or it is only
> > +   * me so stupid? Anathema to whoever designed this non-sense
> > +   * with "intent.open".
> > +   */
> > +  if (!IS_ERR(nd->intent.open.file))
> > +   release_open_intent(nd);
> >    return error;
> > + }
> >   nd->flags &= ~LOOKUP_PARENT;
> >   if (nd->last_type == LAST_BIND)
> >    goto ok;
> >
>
> It's good to have some more Alexeycomments in the tree.
>
> I wonder if we're also needing a path_release() here.  And if not, whether
> it is still safe to run release_open_intent() against this nameidata?
>
> Hopefully Trond can recall what's going on in there...

The patch looks correct, except that I believe we can skip the IS_ERR()
test there: if we're following links then we presumably have not tried
to open any files yet, so the call to release_open_intent(nd) can be
made unconditional.

Cheers,
  Trond