Subject: Re: [PATCH] fdset's leakage
Posted by Andrew Morton on Tue, 11 Jul 2006 09:28:08 GMT
View Forum Message <> Reply to Message

On Tue, 11 Jul 2006 13:05:03 +0400
Kirill Korotaev <dev@openvz.org> wrote:

> Andrew,
>
> > But the code in there is really sick. In all cases we do:
> >
> > free_fdset(foo->open_fds, foo->max_fdset);
> > free_fdset(foo->close_on_exec, foo->max_fdset);
> >
> > How much neater and more reliable would it be to do:
> >
> > free_fdsets(foo);
> >
> > ?
> agree. should I prepare a patch?

Is OK, I'll take care of it later. We want to let your current patch bake as-is in mainline for a while so that we can backport it into 2.6.17.x with more confidence. That's a bit excessive in this case, but the principle is good.

> > Also,
> >
> > nfds = NR_OPEN_DEFAULT;
> > /*
> >   * Expand to the max in easy steps, and keep expanding it until
> >   * we have enough for the requested fd array size.
> >   */
> > do {
> > #if NR_OPEN_DEFAULT < 256
> >   if (nfds < 256)
> >   nfds = 256;
> >   else
> > #endif
> >   if (nfds < (PAGE_SIZE / sizeof(struct file *)))
> >   nfds = PAGE_SIZE / sizeof(struct file *);
> >   else {
> >   nfds = nfds * 2;
> >   if (nfds > NR_OPEN)
> >   nfds = NR_OPEN;
> >   }
> > } while (nfds <= nr);
> >

> >
> > That's going to take a long time to compute if nr > NR_OPEN. I just fixed
> > a similar infinite loop in this function. Methinks this
> >
> > nfds = max(NR_OPEN_DEFAULT, 256);
> > nfds = max(nfds, PAGE_SIZE/sizeof(struct file *));
> > nfds = max(nfds, round_up_pow_of_two(nr + 1));
> > nfds = min(nfds, NR_OPEN);
> >
> > is clearer and less buggy. I _think_ it's also equivalent (as long as
> > NR_OPEN>256). But please check my logic.
> Yeah, I also noticed these nasty loops but was too lazy to bother :)
> Too much crap for my nerves :)
>
> Your logic looks fine for me.

I usually get that stuff wrong.

> Do we have already round_up_pow_of_two() function

yep, in kernel.h.
