
Subject: Re: [PATCH] diskquota: 32bit quota tools on 64bit architectures

Posted by [Mikael Pettersson](#) on Tue, 19 Jun 2007 20:09:06 GMT

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

Andrew Morton writes:

> On Mon, 18 Jun 2007 12:21:47 +0400

> Vasily Tarasov <vtaras@openvz.org> wrote:

>

> > From: Vasily Tarasov <vtaras@openvz.org>

> >

> > This patch should be applied after Arnd Bergmann's patch,

> > that introduces new compat types:

> > <http://lkml.org/lkml/2007/6/15/98>

> >

> > OpenVZ Linux kernel team has discovered the problem

> > with 32bit quota tools working on 64bit architectures.

> > In 2.6.10 kernel sys32_quotactl() function was replaced by sys_quotactl() with

> > the comment "sys_quotactl seems to be 32/64bit clean, enable it for 32bit"

> > However this isn't right. Look at if_dqblk structure:

> >

> > struct if_dqblk {

> > __u64 dqb_bhardlimit;

> > __u64 dqb_bsoftlimit;

> > __u64 dqb_curspace;

> > __u64 dqb_ihardlimit;

> > __u64 dqb_isoftlimit;

> > __u64 dqb_curinodes;

> > __u64 dqb_btime;

> > __u64 dqb_itime;

> > __u32 dqb_valid;

> > };

> >

> > For 32 bit quota tools sizeof(if_dqblk) == 0x44.

> > But for 64 bit kernel its size is 0x48, 'cause of alignment!

> > Thus we got a problem. Attached patch reintroduce sys32_quotactl() function,

> > that handles this and related situations.

> >

> > Signed-off-by: Vasily Tarasov <vtaras@openvz.org>

> >

> > ---

> >

> > In OpenVZ technology 32 bit Virtual Environments over

> > 64 bit OS are common, hence we have customers, that complains on this bad quota

> > behaviour:

> >

> > # /usr/bin/quota

> > quota: error while getting quota from /dev/sda1 for 0: Success

> >

> > The reason is caused above.

>

> Only x86_64 and ia64 are fixed. Would it be correct to assume that the
> other CONFIG_COMPAT architectures also need to be fixed?

I complained about this very issue when a previous version of this patch was submitted last week, and Arnd explained that non-x86 doesn't have a problem here because `alignof(u64)` is the same in 32- and 64-bit modes.

However, the fact that the patch description talks about 32- and 64-bit machines `_in_general_`, while the patch clearly only handles x86-32 on x86-64 and ia64, is itself a bug. A more precise patch description and a better comment in the code is in order, I think.

/Mikael
