
Subject: Re: [patch -mm 2/4] mqueue namespace : add unshare support

Posted by [serue](#) on Thu, 29 Nov 2007 13:57:28 GMT

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

Quoting Andrew Morton (akpm@linux-foundation.org):

> On Thu, 29 Nov 2007 11:28:28 +0100 Cedric Le Goater <clg@fr.ibm.com> wrote:

>

> > > Index: 2.6.24-rc3-mm2/include/linux/sched.h

> > > =====

> > > --- 2.6.24-rc3-mm2.orig/include/linux/sched.h

> > > +++ 2.6.24-rc3-mm2/include/linux/sched.h

> > > @@ -27,6 +27,7 @@

> > > #define CLONE_NEWUSER 0x10000000 /* New user namespace */

> > > #define CLONE_NEWPID 0x20000000 /* New pid namespace */

> > > #define CLONE_NEWNET 0x40000000 /* New network namespace */

> > > +#define CLONE_NEWMQ 0x80000000 /* New posix mqueue namespace */

> > >

> > > That's it :) We've run out of clone flags on 32-bit platforms :(

> >

> > yes.

> >

> > I have been giving some thoughts to a clone2() to extend the flags

>

> There appears to be little alternative.

Just thinking aloud, but

given the concerns with the safety and sanity of unsharing only partial namespaces, and before much userspace is depending on any of

CLONE_NEWUTS,CLONE_NEWIPC,CLONE_NEWUSER,CLONE_NEWNET,CLONE_NEWMQ
UEUE

maybe we should have traditional clone only support CLONE_NEWNS (since it's the most useful on its own) and CLONE_NEWCONTAINER, where CLONE_NEWCONTAINER always unshares all the namespaces we know about.

Then clone2 can allow more finegrained choice of namespaces. It takes the exact same clone_flags as clone(), but instead of parent_tidptr and child_tidptr args it has a ns_unshare flag which specifies which namespaces to unshare.

-serge

> > but

> > andrew is preparing to recycle CLONE_DETACHED and CLONE_STOPPED for

> > 2.6.26. Some we might have some more time in front of us.

>

> CLONE_DETACHED proved to be in use. There are no reports of anyone using

> CLONE_STOPPED though.

>

>

> Containers mailing list

> Containers@lists.linux-foundation.org

> <https://lists.linux-foundation.org/mailman/listinfo/containers>

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
