
Subject: Re: [PATCH 1/2][RFC] containers: improve automatic container naming
Posted by [serue](#) on Fri, 01 Jun 2007 22:38:55 GMT

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

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

> On Fri, 1 Jun 2007 16:48:09 -0500

> "Serge E. Hallyn" <serue@us.ibm.com> wrote:

>

> > This compiles and boots, but is not intended for inclusion in -mm (yet),
> > just as an RFC for the naming scheme to fix the bug Andrew pointed out.

> >

> > Seem ok overall?

> >

> > thanks,

> > -serge

> >

> > >From 8e9b972f7482415777e982d3bc9a0d55cbaf862b Mon Sep 17 00:00:00 2001

> > From: Serge E. Hallyn <serue@us.ibm.com>

> > Date: Fri, 1 Jun 2007 15:32:15 -0400

> > Subject: [PATCH 1/2] containers: improve automatic container naming

> >

> > The automatic naming of containers created using container_clone()
> > is currently broken (not protected from wraparound) and inconvenient.

> >

> > Add a per-container counter for use in naming children of the container.

> > Before two unshares in a row by one process, and a third in another,

> > would result in

> >

> > /node1/node2

> > /node3

> >

> > The current scheme should result in

> >

> > /node1/node1

> > /node2

> >

> > Also, keep a hash table populated with used names, to protect
> > against counter wrap-around.

> >

> > ...

> >

> > include/linux/container.h | 8 +++

> > kernel/container.c | 116 ++++++

>

> gad, what's all this stuff?

The sound of me starting over.

> I think an IDR tree would get you what you're after in much less code.

Ok, will look at that. Thanks for the tip!

> Although it means that container IDs would get recycled quickly across a
> remove+add.

>

> Be aware that there are IDR enhancements in Greg's driver tree (and hence
> in -mm) which are relevant to this application.

>

>

> > + if (cont->auto_cnt_set) {

>

> Can we please stop using "cnt" and "cont" to refer to containers? Let's
> use "container", OK?

Ok.

thanks,
-serge
