
Subject: Re: [PATCH] Filesystems visibility control group

Posted by [serue](#) on Fri, 08 Feb 2008 15:54:48 GMT

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

Quoting Pavel Emelyanov (xemul@openvz.org):

> Dave Hansen wrote:

> > On Thu, 2008-02-07 at 18:04 +0300, Pavel Emelyanov wrote:

> >> Having this proxy is the easiest way to keep the global list and

> >> file_system_type structure (almost) untouched and simplify the code.

> >>

> >> The filesystems.list file syntax is simple: [+<name> without

> >> a '\n' at the end. Made for 2.6.24-rc8-mm1

> >

> > For what would you want this? What kinds of filesystems would you like

> > to restrict the use of in a container?

>

> Any you want to.

>

> E.g. ext3 with its journaling thread. Having 100 containers

> with a journal in each kills the node.

> E.g. cgroupsfs if you don't want your cgroup create its

> own ones.

Ah, so that answers that question - you enforce this by umounting

/cgroups in the container and not letting them remount it?

> E.g. NFS that can stuck on umount and block you container

> on stop.

>

> > -- Dave

>

>

> 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>
