
Subject: Re: [RFC][PATCH 0/7] Clone PTS namespace
Posted by [Sukadev Bhattiprolu](#) on Wed, 09 Apr 2008 16:23:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

H. Peter Anvin [hpa@zytor.com] wrote:

> sukadev@us.ibm.com wrote:

>> Devpts namespace patchset

>> In continuation of the implementation of containers in mainline, we need

>> to

>> support multiple PTY namespaces so that the PTY index (ie the tty names)

>> in

>> one container is independent of the PTY indices of other containers. For

>> instance this would allow each container to have a '/dev/pts/0' PTY and

>> refer to different terminals.

>

> Why do we "need" this? There isn't a fundamental need for this to be a
> dense numberspace (in fact, there are substantial reasons why it's a bad
> idea; the only reason the namespace is dense at the moment is because of
> the hideously bad handling of utmp in glibc.) Other than indices, this
> seems to be a more special case of device isolation across namespaces,
> would that be a more useful problem to solve across the board?

We want to provide isolation between containers, meaning PTYs in container
C1 should not be accessible to processes in C2 (unless C2 is an ancestor).

The other reason for this in the longer term is for checkpoint/restart.

When restarting an application we want to make sure that the PTY indices
it was using is available and isolated.

We started out with isolating just the indices but added the special-case
handling for granting the host visibility into a child-container.

A complete device-namespaces could solve this, but IIUC, is being planned
in the longer term. We are hoping this would provide the isolation in the
near-term without being too intrusive or impeding the implementation of
the device namespaces.

Sukadev

Containers mailing list

Containers@lists.linux-foundation.org

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