Subject: Re: [RFC] [PATCH 0/3] containers: introduction
Posted by ebiederm on Sat, 23 Dec 2006 02:55:10 GMT
View Forum Message <> Reply to Message

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

> Following is a small patchset implementing what I was describing
> about earlier, namely semantics for a hierarchical container
> naming scheme.
>
> What works:
>
> 1. ls -l /proc/$$/container
> shows the full hierarchical name of the container;
>
> 2. mount -t containerfs none /container
> results in a file tree under /container representing the
> full container hierarchy
>
> 3. cd /proc/$$/container; ls
> results in a listing of child containers
>
> What doesn't work:
> 1. The /proc/$$/container link always appears dead (red
> in bash on my fedora test system) because it points
> into a kern_mounted fs.

Just a quick comment. I am not at all comfortable exporting
internal kernel mounts without something explicit happening.
I played with that in one of my earlier patches and the corner cases
are just extremely weird and mess with the usual unix guarantees
about the namespace. Two specific examples are that .. fails to work properly,
as does sys_getcwd.

My gut feel is that we need something like union mounts so we can
glue these kinds of things but that will mount and unmount as
a unit. So we can preserve backwards compatibility with existing
filesystems. I haven't had a chance to look at what it would
take to implement this kind of hidden union mount though.

With the mount tree cloning code I believe we are quite close (at least
if we don't need the union property).

Eric

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers