
Subject: Re: [RFC][PATCH][cryo] Save/restore state of unnamed pipes
Posted by [Dave Hansen](#) on Mon, 23 Jun 2008 23:30:28 GMT

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

On Sun, 2008-06-22 at 17:40 -0400, Oren Laadan wrote:

> One issue with this approach is that the operation (checkpoint) becomes
> non-transparent to the "checkpintee"s: re-writing the data back into
> the pipe/fifo may generate a SIGIO at the read-side of the pipe/fifo, or
> an inotify event on the inode, without a real reason. This may confuse
> the application if its logic relies on "normal" behavior.

Yeah, I guess this is true, somewhat.

But, for something expecting to read from a pipe, they'll get a SIGIO when the data is first written. If the signal has not yet been delivered then we're fine because we won't generate *another* one.

If they've already received the signal and left the data in the pipe, what harm does another SIGIO do? The kernel is free to give the app a SIGIO for every single byte stuck in the pipe, right? So, what difference does any number of SIGIOs make when there's data sitting in the pipe unread?

As for inotify, yeah that might be an actual issue. It's probably good to note this potential issue, and to watch for anyone doing inotifies on a pipe fd when we get to checkpointing inotify.

Overall, seems like an awfully minor issue to me.

-- Dave

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
