

---

Subject: Re: [RFC PATCH 5/5] use next syscall data to predefine the file descriptor value

Posted by [Nadia Derby](#) on Thu, 10 Jul 2008 06:25:48 GMT

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

---

Eric W. Biederman wrote:

> Nadia.Derbey@bull.net writes:

>

>

>>[PATCH 05/05]

>>

>>This patch uses the value written into the next\_syscall\_data proc file

>>as a target file descriptor for the next file to be opened.

>>

>>This makes it easy to restart a process with the same fds as the ones it was

>>using during the checkpoint phase, instead of 1. opening the file, 2. dup2'ing

>>the open file descriptor.

>

>

> As it happens the behavior of open is deterministic. So if you open

> the files in the right order you should not need this. dup2 is only needed

> if there is a gap in the fds used.

>

This covers the case where you're checkpointing a process that has

1. opened, say 3 files (fds x, x+1, and x+2)

2. closed fd x+1

--> checkpoint occurs at that point.

During restart, you'll have to only recreate fds x and x+2.

But I'm realizing that this might be what you're calling a gap in the  
fds ;-)

Regards,  
Nadia

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

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

---