
Subject: Re: [RFC PATCH 0/5] Resend -v2 - Use procfs to change a syscall behavior

Posted by [ebiederm](#) on Thu, 10 Jul 2008 00:43:04 GMT

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

Alexey Dobriyan <adobriyan@gmail.com> writes:

> On Tue, Jul 08, 2008 at 01:24:22PM +0200, Nadia.Derbey@bull.net wrote:
>> # echo "LONG1 XX" > /proc/self/task/<my_tid>/next_syscall_data
>
> Same stuff.
>
> There is struct task_struct::did_exec , what about it?
>
> Also, patches are about de-serializing, how serializing from userspace looks
> like?
> You freezed group of processes, then what?
>
> How, for example, dump all VMAs correctly?
> [prepares counter-example]

Alexey userspace vs a kernel space implementation is the wrong argument.

It is clearly established that the current user space interfaces are insufficient to do the job. So we need to implement something in the kernel.

Further I have heard of no one suggesting running a single kernel on multiple machines. Therefore there no one seems to be doing this entirely in the kernel and so we need a user space component.

So the question should not be user space vs. kernel space but can we build clean interfaces for checkpoint/restart? What will those interfaces be?

Although I think it is good that we are seeing more people play with this as that should mean that our pool of people for doing code review on the implementation should be reasonable.

Eric

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