
Subject: Re: [PATCH 1/2] signal checkpoint: define /proc/pid/sig/
Posted by [Carl-Daniel Hailfinge](#) on Tue, 12 Jun 2007 15:29:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 11.06.2007 19:05, Serge E. Hallyn wrote:

> Quoting Cedric Le Goater (clg@fr.ibm.com):

>

>> should we continue to use /proc ? or switch to some other mechanisms

>> like getnetlink (taskstats) to map kernel structures.

>

> We want to avoid 'map'ping kernel structures, though, right? We can

> dump the data in a more generic fashion through netlink, dunno what we

> prefer. But this is very definately process information :), so /proc

> does seem appropriate.

While I agree that /proc seems appropriate, I see a few benefits of dumping the data through netlink:

- * Speed. IIRC there were benchmarks showing an advantage of netlink over /proc when communicating with userspace. Sorry, no idea where I read that.

- * Versioning. While we strive to have the perfect interface on the first try, changes might be necessary. I see no way to handle multiple versions of an interface in /proc without big headaches.

- * Conformity. With /proc, people often see a file, take a look at it and try to infer the structure of the file from what they see. This has led to multiple problems in the past when the content of some files in /proc changed slightly and tools broke. With netlink, implementers have to look at the spec to achieve anything useful.

Regards,
Carl-Daniel

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