
Subject: Re: [RFC PATCH 3/5] use next syscall data to predefine process ids
Posted by [Nadia Derby](#) on Thu, 10 Jul 2008 08:32:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

Eric W. Biederman wrote:

> Nadia.Derbey@bull.net writes:

>
>
>>[PATCH 03/05]
>>
>>This patch uses the value written into the next_syscall_data proc file
>>as a target upid nr for the next process to be created.
>>The following syscalls have a new behavior if next_syscall_data is set:
>>. fork()
>>. vfork()
>>. clone()
>>
>>In the current version, if the process belongs to nested namespaces, only
>>the upper namespace level upid nr is allowed to be predefined, since there
>>is not yet a way to take a snapshot of upid nrs at all namespaces levels.
>>
>>But this can easily be extended in the future.
>
>
> This patch is unnecessary. The and a mess. The existing limits on the pid range should
> be enough. We may need to export it via /proc/sys.
>

Eric,

If I correctly understood what you're saying, it means set min = max = target_pid using /proc/sys, i.e. for the whole system: don't you think this might be dangerous: allocating pids will fail for any other running process during the entire period of time where /proc/sys will be set like that.

I really think this is a feature that should be confined to a process.

Regards,
Nadia

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