
Subject: Re: [RFC][PATCH] Add child reaper to struct pspace
Posted by [ebiederm](#) on Tue, 12 Sep 2006 15:43:19 GMT
[View Forum Message](#) <> [Reply to Message](#)

Kirill Korotaev <dev@sw.ru> writes:

> I guess there will be a need of list of tasks... not of pids only...
> many of loops like do_each_thread()/while_each_thread() has nothing to do with
> pids
> and should be narrowed down to loop through the container.
>
> Does this logic belong to pid_ns? if yes, then it definitely should be called
> task_ns.

Just skimming through I see one or two instances. Where the existing loop uses do_each_thread()/while_each_thread() that we need to change.

kernel/capabilities.c cap_set_all() is an example.

However what we are trying to achieve there is to iterate through the same list that kill(-1,) uses. So we need to replace do_each_thread()/while_each_thread() with something that will iterate through everything in the pid namespace.

Most instances of do_each_thread()/while_each_thread() the kernel really does need a global view, and need to be left unchanged.

Basically the current kernel is short the concept of a process group of all processes, and uses the concept of a list of all processes instead.

Since the two concepts of a list of all tasks, and a list of all processes I can see diverge when we have multiple pid namespaces we need to add an additional concept, in the kernel.

Do you know an example in that we need to change to implement a pid namespace that goes beyond iterating through the list of processes that kill(-1,) uses?

Eric
