
Subject: Re: [DRAFT] Container mini-summit notes v0.01

Posted by [dev](#) on Thu, 06 Sep 2007 12:00:47 GMT

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

Cedric Le Goater wrote:

```
>>> * possible direction for C/R user api
>>>   . checkpoint/restart syscalls
>>>   . C/R file systems
>>>   solves the set id issue
>>>   elegant but exposes too much the ABI
>>
>>I vote for the filesystem :) I'd add more details as we did on mini-summit.
>>
>>tasks
>>  - <pid1>
>>  - <tid1>
>>    ...
>>    <tidN>
>>    files
>>    1 -> /* made as a symlink */
>>    2 -> /* if socket point to net/ objects */
>>    memory
>>    - <vma1> -> /* symlink to mm objects */
>>  <pid2>
>>    ...
>>  <pidN>
>>mm
>>ipc
>>network
>>
>>and so on and so forth.
>
>
> We need to dig on this idea. RFC ?
```

1. resource interrelationships are much more complicated then a tree.
e.g. pid can be owned by a number of processes, threads, terminals, etc.
So I'm not a fan of the idea.
2. exposing such a low-level information to the user-space can be undesirable:
 - a) it allows to create non-GPL checkpointing
 - b) significantly hits the performance of checkpoint/restore
 - c) BTW, how do you plan to restore via filesystem? mkdir? create? :)

Thanks,
Kirill

Containers mailing list

