

Eric W. Biederman wrote:

>>The overlapping is important if you want to pretend that the  
>>namespace-able resources are allowed to be specified per-process, when  
>>really they are specified per-family.

>>

>>In this way, a process family is merely a grouping of processes with  
>>like namespaces, and depending on which way they overlap you get the  
>>same behaviour as when processes only have one resource different, and  
>>therefore remove the overhead on fork().

>>

>>

>

>

>I missed this subthread originally.

>

>I think it is important that we can have containers in containers  
>if at all possible. This means large software collections can count  
>on them being present.

>

>

Right. Well, my concept was that that "in"-ness is just a relationship between two process families, so treat it relationally and not heirarchically. So, to the kernel, they've all got global unique IDs, but to the actual userspace within those families, they might see something different. And the model still supports containers in containers.

>As for having some items inside a namespace show up in both  
>a parent and a child namespace I think the case is less clearly  
>defined. If possible that is something we want to avoid as it  
>complicates the implementation.

>

>For pids I will be surprised if we can avoid it.

>

>For most other namespaces I think we can, and it is a good thing  
>to avoid.

>

>

Well, let's cross those bridges when we come to them. I agree it should only be implemented if required for individual numbers. PIDs and process family IDs seem logical to me, at this point anyway.

Sam.

---