

---

Subject: Re: [PATCH][DOCUMENTATION] The namespaces compatibility list doc  
Posted by [Pavel Emelianov](#) on Fri, 16 Nov 2007 16:11:25 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Daniel Lezcano wrote:

> Pavel Emelianov wrote:

>> >From time to time people begin discussions about how the  
>> namespaces are working/going-to-work together.

>>

>> Ted T'so proposed to create some document that describes what  
>> problems user may have when he/she creates some new namespace,  
>> but keeps others shared. I liked this idea, so here's the  
>> initial version of such a document with the problems I currently  
>> have in mind and can describe somewhat audibly - the "namespaces  
>> compatibility list".

>>

>> The Documentation/namespaces/ directory is about to contain more  
>> docs about the namespaces stuff.

>>

>> Thanks to Cedirc for notes and spell checks on the doc.

>>

>> Signed-off-by: Pavel Emelianov <xemul@openvz.org>

>>

>> ---

>>

>> commit 83061c56e1c4dcd54d48a62b108d219a7f5279a0

>> Author: Pavel <pavel@xemulnb.sw.ru>

>> Date: Fri Nov 16 12:25:53 2007 +0300

>>

>> Namespaces compatibility list

>>

>> diff --git a/Documentation/00-INDEX b/Documentation/00-INDEX

>> index 910e511..3ead06b 100644

>> --- a/Documentation/00-INDEX

>> +++ b/Documentation/00-INDEX

>> @@ -262,6 +262,8 @@ mtrr.txt

>> - how to use PPro Memory Type Range Registers to increase performance.

>> mutex-design.txt

>> - info on the generic mutex subsystem.

>> +namespaces/

>> + - directory with various information about namespaces

>> nbd.txt

>> - info on a TCP implementation of a network block device.

>> netlabel/

>> diff --git a/Documentation/namespaces/compatibility-list.txt

b/Documentation/namespaces/compatibility-list.txt

>> new file mode 100644

>> index 0000000..9c9e5c1

```

>> --- /dev/null
>> +++ b/Documentation/namespaces/compatibility-list.txt
>> @@ -0,0 +1,33 @@
>> + Namespaces compatibility list
>> +
>> +This document contains the information about the problems user
>> +may have when creating tasks living in different namespaces.
>> +
>> +Here's the summary. This matrix shows the known problems, that
>> +occur when tasks share some namespace (the columns) while living
>> +in different other namespaces (the rows):
>> +
>> + UTS IPC VFS PID User Net
>> +UTS X
>> +IPC X 1
>> +VFS X
>> +PID 1 1 X
>> +User 2 X
>> +Net
>> +
>
> UTS IPC VFS PID User Net
> UTS X
> IPC X 1 3
> VFS X
> PID 1 1 X
> User 2 X
> Net
>
>> +1. Both the IPC and the PID namespaces provide IDs to address
>> + object inside the kernel. E.g. semaphore with ipcid or
>> + process group with pid.
>> +
>> + In both cases, tasks shouldn't try exposing this id to some
>> + other task living in a different namespace via a shared filesystem
>> + or IPC shmem/message. The fact is that this ID is only valid
>> + within the namespace it was obtained in and may refer to some
>> + other object in another namespace.
>> +
>> +2. Intentionnaly, two equal user ids in different user namespaces
> Intentionaly
>> + should not be equal from the VFS point of view. In other
>> + words, user 10 in one user namespace shouldn't have the same
>> + access permissions to files, belonging to user 10 in another
> belonging
>> + namespace. But currently this is not so.
>> +
>

```

> 3. IPC and User

>

> Two processes running in different user namespaces but with the same  
> uids can access with the same permissions to the IPC created by one or  
> other process

Different user and equal IPC is where the row UID and column  
IPC cross, not the vice versa.

Anything, thanks, please send an incremental patch if Andrew  
accepts this one.

Thanks,  
Pavel

---

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

---