

---

Subject: Re: The issues for agreeing on a virtualization/namespaces implementation.

Posted by [ebiederm](#) on Wed, 08 Feb 2006 22:28:12 GMT

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

---

Hubertus Franke <[frankeh@watson.ibm.com](mailto:frankeh@watson.ibm.com)> writes:

> Eric W. Biederman wrote:

>> Hubertus Franke <[frankeh@watson.ibm.com](mailto:frankeh@watson.ibm.com)> writes:

>>

>>> Eric W. Biederman wrote:

>>>

>>

>>>> 3) How do we refer to namespaces and containers when we are not members?

>>>> - Do we refer to them indirectly by processes or other objects that

>>>> we can see and are members?

>>>> - Do we assign some kind of unique id to the containers?

>>>

>>>

>> What I have done which seems easier than creating new names is to refer

>> to the process which has the namespace I want to manipulate.

>

> Is then the idea to only allow the container->init to manipulate

> or is there need to allow other privileged processes to perform namespace

> manipulation?

> Also after thinking about it.. why is there a need to have an external name

> for a namespace ?

There are several cases.

Passing network devices to a child's namespace, as usually the loopback interface is not enough.

Monitoring the namespace from outside, so among other things you aren't required to checkpoint and migrate your monitoring daemon.

There are several other control and monitoring operations that I am not quite as familiar. One of them is the vserver idea of entering a guest.

To expand on things a little bit. If we have interfaces that take strings we can refer to an arbitrary child process as pid/pid/pid/.... So we should not be limited to what is at the init of the container. If that proves desirable.

Permissions checks for most of these operations require some serious thinking before they are merged.

Eric

---