

"Serge E. Hallyn" <serue@us.ibm.com> wrote:

>  
> This patchset introduces a per-process utsname namespace. These can  
> be used by openvz, vserver, and application migration to virtualize and  
> isolate utsname info (i.e. hostname). More resources will follow, until  
> hopefully most or all vserver and openvz functionality can be implemented  
> by controlling resource namespaces from userspace.  
>

Generally, I think that the whole approach of virtualising the OS so it can run multiple independent instances of userspace is a good one. It's an extension and a strengthening of things which Linux is already doing and it pushes further along a path we've been taking for many years. If done right, it's even possible that each of these featurettes could improve the kernel in its own right - better layering, separation, etc.

The approach which you appear to be taking is to separate the bits of functionality apart and to present them as separate works each of which is reviewed-by, acceptable-to and will-be-used-by all of the interested projects. That's ideal, and is very much appreciated.

All of which begs the question "now what?".

What we do not want to do is to merge up a pile of infrastructural stuff which never gets used. On the other hand, we don't want to be in a position where nothing is merged into mainline until the entirety of vserver &&|| openvz is ready to be merged.

I see two ways of justifying a mainline merge of things such as this

a) We make an up-front decision that Linux will have OS-virtualisation capability in the future and just start putting in place the pieces for that, even if some of them are not immediately useful.

I suspect that'd be acceptable, although I worry that we'd get partway through and some issues would come up which are irreconcilable amongst the various groups.

It would help set minds at ease if someone could produce a bullet-point list of what features the kernel will need to get it to the stage where "most or all vserver and openvz functionality can be implemented by controlling resource namespaces from userspace." Then we can discuss that list, make sure that everyone's pretty much in

agreement.

It would be good if that list were to identify which features are useful to Linux in their own right, and which ones only make sense within a whole virtualise-the-OS setup.

- b) Only merge into mainline those feature which make sense in a standalone fashion. eg, we don't merge this patchset unless the "per-process utsname namespace" feature is useful to and usable by a sufficiently broad group of existing Linux users.

I suspect this will be a difficult approach.

The third way would be to buffer it all up in -mm until everything is sufficiently in place and then slam it all in. That might not be feasible for various reasons - please advise..

A fourth way would be for someone over there to run a git tree - you all happily work away, I redistribute it in -mm for testing and one day it's all ready to merge. I don't really like this approach. It ends up meaning that nobody else reviews the new code, nobody else understands what it's doing, etc. It's generally subversive of the way we do things.

Eric, Kirill, Herbert: let us know your thoughts, please.

---