

---

Subject: Re: [PATCH 0/9] namespaces: Introduction

Posted by [serue](#) on Fri, 19 May 2006 13:30:41 GMT

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

---

Quoting Eric W. Biederman (ebiederm@xmission.com):

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

>

> > 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.

> >

> > Previous utsname submissions placed a pointer to the utsname namespace  
> > straight in the task\_struct. This patchset (and the last one) moves  
> > it and the filesystem namespace pointer into struct nsproxy, which is  
> > shared by processes sharing all namespaces. The intent is to keep  
> > the taskstruct smaller as the number of namespaces grows.

>

>

> Previously you mentioned:

> > BTW - a first set of comparison results showed nsproxy to have better  
> > dbench and tbench throughput, and worse kernbench performance. Which  
> > may make sense given that nsproxy results in lower memory usage but  
> > likely increased cache misses due to extra pointer dereference.

>

> Is this still true? Or did our final reference counting tweak fix  
> the kernbench numbers?

I haven't checked that. I'll start a new set of runs later this morning, should get the results out saturday.

-serge

---