

---

Subject: Re: [PATCH 0/9] namespaces: Introduction  
Posted by [ebiederm](#) on Fri, 19 May 2006 08:50:19 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

"Serge E. Hallyn" <[serue@us.ibm.com](mailto: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 just want to be certain that we don't add an optimization,  
that reduces performance.

> Changes:  
> - the reference count on fs namespace and uts namespace now  
> refers to the number of nsproxies pointing to it  
> - some consolidation of namespace cloning and exit code to  
> clean up kernel/{fork,exit}.c  
> - passed ltp and ltpstress on smp power, x86, and x86-64  
> boxes.

Nice.

Eric

---