
Subject: [PATCH 0/9] namespaces: Introduction
Posted by [serue](#) on Thu, 18 May 2006 15:47:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

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.

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.
-