
Subject: Re: [PATCH 4/9] namespaces: utsname: switch to using uts namespaces
Posted by [ebiederm](#) on Fri, 19 May 2006 11:58:22 GMT

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

"Randy.Dunlap" <rdunlap@xenotime.net> writes:

> OK, here's my big comment/question. I want to see <nodename> increased to
> 256 bytes (per current POSIX), so each field of struct <variant>_utsname
> needs be copied individually (I think) instead of doing a single
> struct copy.
>
> I've been working on this for the past few weeks (among other
> things). Sorry about the timing.
> I could send patches for this against mainline in a few days,
> but I'll be glad to listen to how it would be easiest for all of us
> to handle.
>
> I'm probably a little over half done with my patches.
> They will end up adding a lib/utsname.c that has functions for:
> put_oldold_unmame() // to user
> put_old_uname() // to user
> put_new_uname() // to user
> put_posix_uname() // to user

Looking 256 at least makes sense to hold a dns fully qualified domain
name. So even if it isn't specified by posix it make sense.

Can we please make the structure we return to user space look something
like:

```
struct long_utsname {  
    char *sysname;  
    char *nodename;  
    char *release;  
    char *version;  
    char *machine;  
        char *domainname;  
        char buf[0];  
}
```

```
int sys_long_uname(char *buf, size_t bufsz);
```

So we don't hard code the maximum length of these strings into the user
interface, and can just return more by increasing our buffer size.

Eric
