Subject: [RFC][PATCH 0/5] uts namespaces: Introduction
Posted by serue on Fri, 07 Apr 2006 18:36:00 GMT
View Forum Message <> Reply to Message

Introduce utsname namespaces. Instead of a single system_utsname containing hostname domainname etc, a process can request it's copy of the uts info to be cloned. The data will be copied from it's original, but any further changes will not be seen by processes which are not it's children, and vice versa.

This is useful, for instance, for vserver/openvz, which can now clone a new uts namespace for each new virtual server.

This patchset is based on Kirill Korotaev's Mar 24 submission, taking comments (in particular from James Morris and Eric Biederman) into account.

Some performance results are attached. I was mainly curious whether it would be worth putting the task_struct->uts_ns pointer inside a #ifdef CONFIG_UTS_NS. The result show that leaving it in when CONFIG_UTS_NS=n has negligable performance impact, so that is the approach this patch takes.

-serge

Performance testing was done on a 2-cpu hyperthreaded
x86 box with 16G ram. The following tests were run:
dbench (20 times, four clients, on reiser fs non-isolated partition)
tbench (20 times, 5 connections)
kernbench (20 times)
reaim (20 times ranging from 1 to 15 users)

They were run on 2.6.17-rc1:
pristine
patched, but with !CONFIG_UTS_NS ("disabled")
patched with CONFIG_UTS_NS=y ("enabled")

All results are presented as means +/- 95% confidence interval.

Dbench results:
pristine: 387.080727 +/- 9.344585
patched disabled: 389.524364 +/- 9.574921
patched enabled: 370.155600 +/- 30.127808

Tbench results:
pristine: 388.940100 +/- 18.095104
patched disabled: 389.173700 +/- 23.658035
patched enabled: 394.333200 +/- 25.813393

Kernbench results:
pristine: 70.317500 +/- 0.210833
patched, disabled: 70.860000 +/- 0.179292
patched, enabled: 70.346500 +/- 0.184784

Reaim results:
pristine:

Nclients	Mean	95% CI
1	106080.000000	11327.896029
3	236057.142000	18205.544810
5	247867.136000	23536.800062
7	265370.000000	21284.335743
9	262969.936000	18225.497529
11	278256.000000	6230.342816
13	284288.016000	8924.589388
15	286987.170000	7881.034658

patched, disabled:

Nclients	Mean	95% CI
1	105400.000000	8739.978241
3	229500.000000	0.000000
5	252325.176667	16685.663423
7	265125.000000	6747.777319
9	271258.645000	11715.635212
11	280662.608333	7775.229351
13	277719.706667	8173.390359
15	278515.421667	10963.211450

patched, enabled:

Nclients	Mean	95% CI
1	102000.000000	0.000000
3	224400.000000	14159.870036
5	242963.288000	40529.490781
7	255150.000000	8745.802081
9	270154.284000	8918.863136
11	283134.260000	12239.361252
13	288497.540000	11336.550964
15	280022.728000	8804.882369