Subject: [PATCH 0/28] Pid namespaces (two models)
Posted by Pavel Emelianov on Fri, 15 Jun 2007 15:55:43 GMT
View Forum Message <> Reply to Message

Long ago Sukadev and I sent two approaches for pid namespaces - the
hierarchical model in which namespaces are nested into each other,
and the flat model, where pids have only two values and creation of
level 3 namespace is prohibited.

After that I showed that multilevel model introduces a noticeable
overhead of approximately 1-2% to kernel standard operations like
fork() and getpid(). At the same time flat model showed no performance
hit on these tests.

Nevertheless multilevel model is worth living.

This set introduces booth models each under its config option. The
set is logically splitted into the following parts:
* [PREP] subset - the preparations for the namespaces. These patches
  by their own do not change the kernel behavior, but prepare the ground
  for the pid namespaces. This subset weights 14 patches;
* [FLAT] subset - the flat namespaces model. This is 6 patches;
* [MULTI] subset - the multilevel models. This is 6 patches also;
* Patch for cloning the namespace;
* Patch with Kconfig options.

The following tests were run:
[1] nptl perf test
[2] getpid() speed
[3] ltp (not for speed, but for kernel API checks)

The testing results summary:
* Flat model provides zero overhead in init namespace for all the
tests and less than 7% in the namespace for nptl test only.
* Multilevel model provides up to 2% overhead in init namespace and
more than 10% for nptl test in the level 2 namespace.

Testing details:

	perf, s	getpid
after unshare		

--------------------+-------------+-----------+

[1] ./perf -s 1000000 -t 1 -r 0 -T --sync-join
    nptl/perf.c from glibc-2.5

    error is 3 standard deviations

[2] getpid(2) done 10^9 times real time as reported by time(1)

The patches are for 2.6.22-rc4-mm2 tree.

Thanks,
Pavel

_____
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers