
Subject: Re: [RFC][PATCH 0/15] Pid namespaces
Posted by [Sukadev Bhattiprolu](#) on Fri, 27 Jul 2007 04:22:13 GMT
[View Forum Message](#) <> [Reply to Message](#)

Pavel,

We seem to have a memory leak. Its new in this patchset (i.e the following test ran fine on the 2.6.22-rc6-mm1 patchset).

To repro: run "pidns_exec ./mypid" in a tight-loop - where mypid.c is:

```
#include <stdio.h>
#include <unistd.h>

main()
{
    printf("Pid %d, Ppid %d, Pgid %d, Sid %d\n",
           getpid(), getppid(), getpgid(0), getsid(0));
}
```

I ran into OOM in about 30 mins. I am still investigating.

BTW, can we include a simple test program like the pidns_exec in this patch-0 for whoever want to play with pidns ?

Suka

Pavel Emelianov [xemul@openvz.org] wrote:

| (Comment is taken from Sukadev's patchset-v3)

| A pid namespace is a "view" of a particular set of tasks on the system.
| They work in a similar way to filesystem namespaces. A file (or a process)
| can be accessed in multiple namespaces, but it may have a different name
| in each. In a filesystem, this name might be /etc/passwd in one namespace,
| but /chroot/etc/passwd in another.

| For processes, a process may have pid 1234 in one namespace, but be pid 1
| in another. This allows new pid namespaces to have basically arbitrary
| pids, and not have to worry about what pids exist in other namespaces.
| This is essential for checkpoint/restart where a restarted process's pid
| might collide with an existing process on the system's pid.

| In this particular implementation, pid namespaces have a parent-child
| relationship, just like processes. A process in a pid namespace may see
| all of the processes in the same namespace, as well as all of the processes
| in all of the namespaces which are children of its namespace. Processes may
| not, however, see others which are in their parent's namespace, but not in

| their own. The same goes for sibling namespaces.

| This set is based on my patches, I sent before, but it includes some
| comments
| and patches that I received from Sukadev. Sukadev, please, add your
| Acked-by,
| Signed-off-by or From, to patches you want (everybody is also welcome :)).

| The set is based on 2.6.23-rc1-mm1, which already has some preparation
| patches
| for pid namespaces. After the review and fixing all the comments, this set
| will be benchmarked and sent to Andrew for inclusion in -mm tree.

| Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

| Signed-off-by: Sukadev Bhattiprolu <sukadev@us.ibm.com>
