
Subject: Re: [PATCH 3/3] Dynamic kmem cache allocator for pid namespaces

Posted by [Sukadev Bhattiprolu](#) on Sat, 14 Jul 2007 01:22:25 GMT

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

Pavel Emelianov [xemul@openvz.org] wrote:

| Add kmem_cache to pid_namespace to allocate pids from.

| Since booth implementations expand the struct pid to carry
| more numerical values each namespace should have separate
| cache to store pids of different sizes.

| Each kmem cache is names "pid_<NR>", where <NR> is the number
| of numerical ids on the pid. Different namespaces with same
| level of nesting will have same caches.

| This patch has two FIXMEs that are to be fixed after we reach
| the consensus about the struct pid itself.

| The first one is that the namespace to free the pid from in
| free_pid() must be taken from pid. Now the init_pid_ns is
| used.

| The second FIXME is about the cache allocation. When we do know
| how long the object will be then we'll have to calculate this
| size in create_pid_cachep. Right now the sizeof(struct pid)
| value is used.

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

| Acked-by: Cedric Le Goater <clg@fr.ibm.com>

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

| ---

| diff --git a/include/linux/pid_namespace.h b/include/linux/pid_namespace.h

| index ddb9a4c..27cfad3 100644

| --- a/include/linux/pid_namespace.h

| +++ b/include/linux/pid_namespace.h

| @@ -20,6 +20,7 @@ struct pid_namespace {

| struct pidmap pidmap[PIDMAP_ENTRIES];

| int last_pid;

| struct task_struct *child_reaper;

| + struct kmem_cache_t *pid_cachep;

Shouldn't this be 'struct kmem_cache *' ?

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
