

---

Subject: Re: [PATCH 4/7] Implement get\_pts\_ns() and put\_pts\_ns()  
Posted by [Sukadev Bhattiprolu](#) on Tue, 25 Mar 2008 18:44:16 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Serge E. Hallyn [serue@us.ibm.com] wrote:

| Quoting sukadev@us.ibm.com (sukadev@us.ibm.com):

| >  
| > From: Sukadev Bhattiprolu <sukadev@us.ibm.com>  
| > Subject: [PATCH 4/7]: Implement get\_pts\_ns() and put\_pts\_ns()  
| >  
| > Implement get\_pts\_ns() and put\_pts\_ns() interfaces.  
| >  
| > Signed-off-by: Sukadev Bhattiprolu <sukadev@us.ibm.com>  
| > ---  
| > include/linux/devpts\_fs.h | 21 ++++++++  
| > 1 file changed, 20 insertions(+), 1 deletion(-)  
| >  
| > Index: 2.6.25-rc5-mm1/include/linux/devpts\_fs.h  
| > =====  
| > --- 2.6.25-rc5-mm1.orig/include/linux/devpts\_fs.h 2008-03-24 20:04:31.000000000 -0700  
| > +++ 2.6.25-rc5-mm1/include/linux/devpts\_fs.h 2008-03-24 20:05:05.000000000 -0700  
| > @@ -27,13 +27,26 @@ struct pts\_namespace {  
| > extern struct pts\_namespace init\_pts\_ns;  
| >  
| > #ifdef CONFIG\_UNIX98\_PTYS  
| > -  
| > int devpts\_new\_index(void);  
| > void devpts\_kill\_index(int idx);  
| > int devpts\_pty\_new(struct tty\_struct \*tty); /\* mknod in devpts \*/  
| > struct tty\_struct \*devpts\_get\_tty(int number); /\* get tty structure \*/  
| > void devpts\_pty\_kill(int number); /\* unlink \*/  
| >  
| > +static inline void free\_pts\_ns(struct kref \*ns\_kref) { }  
| > +  
| > +static inline struct pts\_namespace \*get\_pts\_ns(struct pts\_namespace \*ns)  
| > +{  
| > + if (ns && (ns != &init\_pts\_ns))  
| > + kref\_get(&ns->kref);  
| > + return ns;  
| > +}  
| > +static inline void put\_pts\_ns(struct pts\_namespace \*ns)  
| > +{  
| > + if (ns && (ns != &init\_pts\_ns))  
| > + kref\_put(&ns->kref, free\_pts\_ns);  
| >

| This isn't right, or I'm not thinking right. Don't you somewhere  
| need to  
|

| 1. rcu\_assign ns->mnt->mnt\_sb->s\_fs\_info to NULL  
| 2. wait a grace period  
| 3. call free\_pts\_ns and check the refcount on the ns again?  
|  
| and then do pts\_ns\_from\_inode() under an rcu\_read\_lock and grab  
| a ref to the ns?

Yes, we need the rcu to grab the reference to pts\_ns.

---

Containers mailing list

Containers@lists.linux-foundation.org

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

---