Subject: Re: [IPC]: Logical refcount loop in ipc ns -> massive leakage
Posted by Alexey Kuznetsov on Mon, 05 Feb 2007 10:14:02 GMT
View Forum Message <> Reply to Message

Hello!

> The struct file that is used appears impossible for user space
> to get at directly. Therefore I believe we can instead increment
> and decrement the namespace count at the same places we increment
> and decrement shm_nattach. Ideally we would only increment the
> namespace count when shm_nattach goes from 0 to 1 and we would
> only decrement the namespace count when shm_nattach goes from 1 to 0.
>
> Does that make sense?

Yes, this would save the day.

Indeed, shm_file_ns() is required only when the segment is already mapped,
except for shm_mmap() and even there shm_nattch is incremented before
do_mmap() is used. It will work.

Possibility to use this file directly will be lost. It is a little unpleasant;
openvz checkpointing used it to restore sysv shm mappings like another file
mappings, it was nice, but this code can be a little uglified to treat
those mapping specially. No harm either.

Alexey

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers