
Subject: Re: [RFC][PATCH 2/2] Virtualization of IPC
Posted by [ebiederm](#) on Fri, 24 Mar 2006 20:09:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

Kirill Korotaev <dev@sw.ru> writes:

> This patch introduces IPC namespaces, which allow to create isolated IPC users
> or containers.
> Introduces CONFIG_IPC_NS and ipc_namespace structure.
> It also uses current->ipc_ns as a pointer to current namespace, which reduces
> places where additional argument to functions should be added.

I don't see where we are freeing the shared memory segments,
the message queues and the semaphores when the last user of the namespace
goes away. Am I missing something?

```
> --- a/include/linux/ipc.h
> +++ b/include/linux/ipc.h
> @@ -70,6 +70,50 @@ struct kern_ipc_perm
>
> #endif /* __KERNEL__ */
>
> +#include <linux/config.h>
> +
> +#ifdef CONFIG_IPC_NS
> +#include <asm/atomic.h>
> +
> +struct ipc_ids;
> +struct ipc_namespace {
> + atomic_t cnt;
> +
> + struct ipc_ids *sem_ids;
> + int sem_ctls[4];
> + int used_sems;
> +
> + struct ipc_ids *msg_ids;
> + int msg_ctlmax;
> + int msg_ctlmnb;
> + int msg_ctlmni;
> +
> + struct ipc_ids *shm_ids;
> + size_t shm_ctlmax;
> + size_t shm_ctlall;
> + int shm_ctlmni;
> + int shm_total;
> +};
```

I believe there is a small problem with this implementation.

per namespace counts and limits are fine. But I think we want to maintain true global limits as well. I know concerns of that nature have been expressed in regards to Daves patch.

```
> --- a/kernel/fork.c
> +++ b/kernel/fork.c
> @@ -1193,6 +1193,7 @@ static task_t *copy_process(unsigned lon
>  attach_pid(p, PIDTYPE_TGID, p->tgid);
>  attach_pid(p, PIDTYPE_PID, p->pid);
>  get_uts_ns(p->uts_ns);
> + get_ipc_ns(p->ipc_ns);
>
>  nr_threads++;
>  total_forks++;
```

Again please move the get outside of the tasklist_lock.

Eric
