
Subject: Re: [IPC]: Logical refcount loop in ipc ns -> massive leakage
Posted by [ebiederm](#) on Sat, 03 Feb 2007 02:47:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

Kirill Korotaev <dev@sw.ru> writes:

> Guys,
>
> Though I have no patch in the hands for mainstream,
> I feel a responsibility to report one majore problem
> related to IPC namespace design.
>
> The problem is about refcounting scheme which is used.
> There is a leak in IPC namespace due to refcounting loop:
> shm segment holds a file, file holds namespace,
> namespace holds shm segment. Loop.
> I suppose the problem is not only IPC-related
> and will happen with some other namespaces as well so should
> be a good lesson for us.
>
> The question is how to fix this.
>
> In OpenVZ we always used 2 different refcounters exactly for this purposes:
> process counter and reference counter.
> When the process counter becomes zero (i.e. the last process from the
> namespace dies) namespace objects are destroyed and cleaned up.
> And the reference counter on the namespace as always protects the structure
> memory only.
>
> How to fix this in mainstream?
> Sure the same approach as above can be used. However, I dislike
> the idea of adding process-counter to each namespace requiring this.
> Any ideas?

I'm slowly beginning to digest this, I don't quite follow what the loop really is yet.

If we don't get to the point where we need multiple counters process counter's are not quite the right concept. We need counters from things that keep the namespace alive.

An open file descriptor to a shm segment needs to keep the namespace alive.

A process attached to the ipc namespace needs to keep the namespace alive.

I will have to look at the code closely to see how what you are

describing can occur, and what we can do to preserve the previous two properties.

Eric

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>
