Subject: [IPC]: Logical refcount loop in ipc ns -> massive leakage
Posted by dev on Wed, 31 Jan 2007 16:48:49 GMT

View Forum Message <> Reply to Message

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 cleanuped.
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?

The relevant OpenVZ patch can be found here:
http://git.openvz.org/?p=linux-2.6.18-openvz;a=commit;h=b11c6ed6e92f0f2291217751596d7d764
6b3ea17

Thanks,
Kirill

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