
Subject: Re: [RFC][PATCH 0/3] Kernel memory accounting container (v2)
Posted by [KAMEZAWA Hiroyuki](#) on Thu, 13 Sep 2007 10:19:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, 13 Sep 2007 13:11:35 +0400
Pavel Emelyanov <xemul@openvz.org> wrote:

> First of all - why do we need this kind of control. The major
> "pros" is that kernel memory control protects the system
> from DoS attacks by processes that live in container. As our
> experience shows many exploits simply do not work in the
> container with limited kernel memory.
>
> I can split the kernel memory container into 4 parts:
>
> 1. kmalloc-ed objects control
> 2. vmalloc-ed objects control
> 3. buddy allocated pages control
> 4. kmem_cache_alloc-ed objects control
>
<snip>
> To play with it, one need to mount the container file system
> with -o kmem and then mark some caches as accountable via
> /sys/slab/<cache_name>/cache_account.
>
Hmm, how can we know "How many kmem will we need ?" in precise per-object
style ? Is this useful ?

Following kind of limitation of user friendly params is bad ?

- # of file handles
- # of tasks
- # of sockets/ connections / packets
- # of posix IPC related things
- and other sources of DoS.

Thanks,
-Kame

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
