
Subject: Re: [PATCH 1/4] Virtualization/containers: introduction
Posted by [Jeff Dike](#) on Thu, 09 Feb 2006 17:47:58 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, Feb 09, 2006 at 11:38:31AM -0500, Hubertus Franke wrote:
> Jeff, interesting, but won't that post some serious scalability issue?
> Imaging 100s of container/namespace ?

In terms of memory?

Running size on sched.o gives me this on x86_64:
text data bss dec hex filename
35685 6880 28800 71365 116c5 sched.o

and on i386 (actually UML/i386)

```
text data bss dec hex filename
10010 36 2504 12550 3106 obj/kernel/sched.o
```

I'm not sure why there's such a big difference, but 100 instances adds a meg or two (or three) to the kernel. This overstates things a bit because there are things in sched.c which wouldn't be duplicated, like the system calls.

How big a deal is that on a system which you plan to have 100s of containers on anyway?

It's heavier than your namespaces, but does have the advantage that it imposes no cost when it's not being used.

> The namespace is mainly there to identify which data needs to be private
> when multiple instances of a subsystem are considered and
> encapsulate that data in an object/datastructure !

Sure, and that's a fine approach. It's just not the only one.

Jeff
