
Subject: Re: Container Test Campaign

Posted by [Cedric Le Goater](#) on Thu, 22 Jun 2006 21:51:10 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi marc !

Marc E. Fiuczynski wrote:

> You mention that testing isolation properties is more of an extra than an
> immediate criteria. Based on our experience, this actually is a fairly
> important criteria. Without decent isolation (both from a namespace and
> resource perspective) it is rather difficult to support lots of concurrent
> users. As our paper states, we run anywhere from 30-90 vservers per machine
> (each machine usually with a 2GHz processor and 1GB of RAM).

is that a common setup for planet lab or a maximum ? how many vservers/
vcontext do you think we should try to reach ?

> We are interested in checkpoint/restart too, but have nothing to test /
> contribute. I've forwarded your message to Jason Nieh @ Columbia. He has a
> relatively long history of working in that area. I saw a demo of their
> checkpoint/restart/migration support last December (live video migrated
> between servers within a single IBM blade system).

we've worked a few years with a zap guy. I only wished they were bit more
open (source) about what they've been doing since crak.

> Their latest paper
> published at USENIX LISA also states that they can migrate from one linux
> kernel version to another. This enables "live" system upgrade, which IMHO
> is just as important as load balancing.

this feature is one the *major* features of mobile containers but it will
require specific kernel APIs to make it maintainable on the long term.

> Another area we are quite interested in is "network virtualization" (private
> route tables, ip tables, etc). We are aware that other container based
> systems (e.g., openvz) have support for this, but we (i.e., PlanetLab) are
> pretty much a vserver shop at the moment. We added our own support to
> safely share a single, public IPv4 address between multiple containers,
> while simultaneously support raw sockets etc. This is an absolute
> requirement for PlanetLab, and I'd argue (but not here) that it also is
> important for desktop usage scenarios that involve containers and want to
> avoid the use of NAT.

Did you contribute that feature to vserver ?

So you have different containers exposing the same IP address ? How do you

assign incoming packets to a container ?

thanks,

C.
