
Subject: Re: [Vserver] Re: Container Test Campaign
Posted by [Herbert Poetzl](#) on Wed, 05 Jul 2006 12:13:40 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, Jul 04, 2006 at 06:32:04PM +0400, Kirill Korotaev wrote:

>>> from the tests:

>>> "For benches inside real 'guest' nodes (OpenVZ/VServer) you should
>>> take into account that the FS tested is not the 'host' node one's."

>>>

>>> at least for Linux-VServer it should not be hard to avoid the
>>> chroot/filesystem namespace part and have it run on the host fs.
>>> a bind mount into the guest might do the trick too, if you need
>>> help to accomplish that, just let me know ...

>> For the moment I just use the chcontext command to get rid of
>> filesystem part. But even if the tested filesystem is not the host
>> filesystem, I just keep in mind that all applications running inside
>> a 'guest' will use `_this_` filesystem and not the host one.

>> From what I understand about VServer, it looks flexible enough to let
>> us test different 'virtualisation' parts. A 'guest' looks like a stack
>> of different 'virtualisation' layers (chcontext + ipv4root + chroot).
>> But it's not the case for all solutions.

> For OpenVZ it is also possible to test different subsystems separately
> (virtualization/isolation, resource management, disk quota, CPU
> scheduler).

> I would notice also, that in OpenVZ all these features are ON by
> default.

which is the same as for a complete guest setup, what I
think (and you already mentioned too, IIRC) is that it is
very important to have identical test setups with and
without virtualization enabled, which means that the
following conditions are met:

- filesystem and involved devices are identical
- used resources and limits are identical/very close
- number of processes and cache state as close as possible
- kernel/system state as close as possible

> So I probably miss something, but why we test other technologies in
> modes when not all the features are ON?

doesn't make much sense, and is not what I actually
suggested in the first place, it might be interesting
to narrow down possible issues by putting together

a complete 'stack' slice by slice if that allows to
remove a single slice from that equation

> in this case we compare not the real overhead, but the one minimized
> for this concrete benchmark.

actually all cases are interesting as none of them
is supposed to add measurable overhead which is also
true for the whole thing which is at least the sum
of all of them

HTC,
Herbert

> It's just like comparing with Xen Domain0 which doesn't have any
> overhead, but not because it is a good technology, but rather because
> it doesn't do anything valuable.
>
> BTW, comparing with Xen would be interesting as well. Just to show the
> difference.
>
> Thanks,
> Kirill
