
Subject: Re: Re: [RFC] Virtualization steps
Posted by [Herbert Poetzl](#) on Thu, 13 Apr 2006 22:45:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, Apr 13, 2006 at 11:33:13PM +0200, Cedric Le Goater wrote:

> Herbert Poetzl wrote:

>

> > well, if the 'results' and 'methods' will be made

> > public, I can, until now all I got was something

> > along the lines:

> >

> > "Linux-VServer is not stable! WE (swsoft?) have

> > a secret but essential test suite running two

> > weeks to confirm that OUR kernels ARE stable,

> > and Linux-VServer will never pass those tests,

> > but of course, we can't tell you what kind of

> > tests or what results we got"

> >

> > which doesn't help me anything and which, to be

> > honest, does not sound very friendly either ...

>

> Recently, we've been running tests and benchmarks in different

> virtualization environments : openvz, vserver, vserver in a minimal

> context and also Xen as a reference in the virtual machine world.

>

> We ran the usual benchmarks, dbench, tbench, lmbench, kernerl build,

> on the native kernel, on the patched kernel and in each virtualized

> environment. We also did some scalability tests to see how each

> solution behaved. And finally, some tests on live migration. We didn't

> do much on network nor on resource management behavior.

I would be really interested in getting comparisons
between vanilla kernels and linux-vserver patched
versions, especially vs2.1.1 and vs2.0.2 on the
same test setup with a minimum difference in config

I doubt that you can really compare across the
existing virtualization technologies, as it really
depends on the setup and hardware

> We'd like to continue in an open way. But first, we want to make sure
> we have the right tests, benchmarks, tools, versions, configuration,
> tuning, etc, before publishing any results :) We have some materials
> already but before proposing we would like to have your comments and
> advices on what we should or shouldn't use.

In my experience it is extremely hard to do 'proper'
comparisons, because the slightest change of the

environment can cause big differences ...

here as example, a kernel build (-j99) on 2.6.16
on a test host, with and without a chroot:

without:

```
451.03user 26.27system 2:00.38elapsed 396%CPU
449.39user 26.21system 1:59.95elapsed 396%CPU
447.40user 25.86system 1:59.79elapsed 395%CPU
```

now with:

```
490.77user 24.45system 2:13.35elapsed 386%CPU
489.69user 24.50system 2:12.60elapsed 387%CPU
490.41user 24.99system 2:12.22elapsed 389%CPU
```

now is chroot() that imperformant? no, but the change
in /tmp being on a partition vs. tmpfs makes quite
some difference here

even moving from one partition to another will give
measurable difference here, all within a small margin

an interesting aspect is the gain (or loss) you have
when you start several guests basically doing the
same thing (and sharing the same files, etc)

> Thanks for doing such a great job on lightweight containers,

you're welcome!

best,
Herbert

> C.
