Subject: Re: Re: [RFC] Virtualization steps
Posted by Cedric Le Goater on Fri, 14 Apr 2006 09:56:21 GMT

View Forum Message <> Reply to Message

Bonjour !

Herbert Poetzl wrote:

> 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

We did the tests last month and used the stable version : vs2.0.2rc9 on a 2.6.15.4. Using benchmarks like dbench, tbench, lmbench, the vserver patch has no impact, vserver overhead in a context is hardly measurable (<3%), same results for a debian sarge running in a vserver.

It is pretty difficult to follow everyone patches. This makes the comparisons difficult so we chose to normalize all the results with the native kernel results. But in a way, this is good because the goal of these tests isn't to compare technologies but to measure their overhead and stability. And at the end, we don't care if openvz is faster than vserver, we want containers in the linux kernel to be fast and stable, one day :)

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

I agree these are very different technologies but from a user point of view, they provide a similar service. So, it is interesting to see what are the drawbacks and the benefits of each solution. You want fault containment and strict isolation, here's the price. You want performance, here's another.

Anyway, there's already enough focus on the virtual machines so we should focus only on lightweight containers.

>> 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

very interesting thanks.

> 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)

we have these in the pipe also, we called them scalability test: trying to
run as much containers as possible and see how performance drops (when the
kernel survives the test :)

ok, now i guess we want to make some kind of test plan.

C.