Subject: Re: Re: [RFC] Virtualization steps
Posted by kir on Tue, 28 Mar 2006 20:50:05 GMT

View Forum Message <> Reply to Message

Jun OKAJIMA wrote:

```
>>I'll summarize it this way: low-level virtualization uses resource
>>inefficiently.
>>
>>With this higher-level stuff, you get to share all of the Linux caching,
>>and can do things like sharing libraries pretty naturally.
>>
>>They are also much lighter-weight to create and destroy than full
>>virtual machines.  We were planning on doing some performance
>>comparisons versus some hypervisors like Xen and the ppc64 one to show
>>scaling with the number of virtualized instances.  Creating 100 of these
>>Linux containers is as easy as a couple of shell scripts, but we still
>>can't find anybody crazy enough to go create 100 Xen VMs.
>>
>>Anyway, those are the things that came to my mind first.  I'm sure the
>>others involved have their own motivations.
>>
>>
>>
>
>Some questions.
>
>1. Your point is rignt in some ways, and I agree with you.
>   Yes, I currently guess Jail is quite practical than Xen.
>   Xen sounds cool, but really practical? I doubt a bit.
>   But it would be a narrow thought, maybe.
>   How you estimate feature improvement of memory shareing
>   on VM ( e.g. Xen/VMware)?
>   I have seen there are many papers about this issue.
>   If once memory sharing gets much efficient, Xen possibly wins.
>
>
```

This is not just about memory sharing. Dynamic resource management is hardly possible in a model where you have multiple kernels running; all of those kernel were designed to run on a dedicated hardware. As it was pointed out, adding/removing memory from a Xen guest during runtime is tricky.

Finally, multiple-kernels-on-top-of-hypervisor architecture is just more complex and has more overhead then one-kernel-with-many-namespaces.

```
>2. Folks, how you think about other good points of Xen,
>   like live migration, or runs solaris, or has suspend/resume or...
```

>
>
OpenVZ will have live zero downtime migration and suspend/resume some time next month.

>   No Linux jails have such feature for now, although I dont think
>   it is impossible with jail.
>
>
>My current suggestion is,
>
>1. Dont use Xen for running multiple VMs.
>2. Use Xen for better admin/operation/deploy... tools.
>
>
This point is controversial. Tools are tools -- they can be made to support Xen, Linux VServer, UML, OpenVZ, VMware -- or even all of them!

But anyway, speaking of tools and better admin operations, what it takes to create a Xen domain (I mean create all those files needed to run a new Xen domain), and how much time it takes? Say, in OpenVZ creation of a VE (Virtual Environment) is a matter of unpacking a ~100MB tarball and copying 1K config file -- which essentially means one can create a VE in a minute. Linux-VServer should be pretty much the same.

Another concern is, yes, manageability. In OpenVZ model the host system can easily access all the VPSs' files, making, say, a mass software update a reality. You can easily change some settings in 100+ VEs very easy. In systems based on Xen and, say, VMware one should log in into each system, one by one, to administer them, which is not unlike the 'separate physical server' model.

>3. If you need multiple VMs, use jail on Xen.
>
>
Indeed, a mixed approach is very interesting. You can run OpenVZ or Linux-VServer in a Xen domain, that makes a lot of sense.