Subject: Xen vs. OpenVZ: live migration
Posted by cdevidal on Wed, 07 Jun 2006 15:15:04 GMT
View Forum Message <> Reply to Message

I'm trying to decide between Xen and OpenVZ. Live migration is quite the "killer app" and now that OpenVZ supports it I have some questions...

* How fast (your best guess) is it over a gigabit network with a typical guest distro? Xen's live migration downtime is microseconds; when I read that OpenVZ has "zero downtime" migration, does that mean it's microseconds, too?
* How does it work without shared hardware? Rsync+SSH?
* What if I had shared hardware, would I also need a lock manager, fencing, etc?

Subject: Re: Xen vs. OpenVZ: live migration
Posted by kir on Wed, 07 Jun 2006 15:27:04 GMT
View Forum Message <> Reply to Message

Quote:I'm trying to decide between Xen and OpenVZ. Live migration is quite the "killer app" and now that OpenVZ supports it I have some questions...
IMHO although live migration is the killer feature this is not the only feature of virtualization, and this is definitely not the feature you use every day. So, to decide which technology is best for you, you need to do a complex evaluation of both Xen and OpenVZ, find out their strengths and weaknesses.

To my mind, the good side of Xen is ability to run different kernels and different OSs (Linux side-by-side to NetBSD, for example). So if you really want to use the different kernels on the same piece of hardware, Xen is your choice.

From the other side, OpenVZ's single kernel approach results in better density, scalability, ability to do mass management and dynamically change all the resources.

Also, search this forum for Xen to get an idea of what users are saying on the Xen vs. OpenVZ subject.

Quote:* How fast (your best guess) is it over a gigabit network with a typical guest distro? Xen's live migration downtime is microseconds; when I read that OpenVZ has "zero downtime" migration, does that mean it's microseconds, too?

This is not really zero downtime -- it is zero downtime from the user's point of view, since user does not see it as a downtime. I guess the delay in processing is a few seconds (non-shared storage scenario). We are improving this time.

Quote:* How does it work without shared hardware? Rsync+SSH?
Exactly

Quote:* What if I had shared hardware, would I also need a lock manager, fencing, etc?
If you mean shared storage (SAN) -- I do not have any experience with this.

Subject: Re: Xen vs. OpenVZ: live migration
Posted by cdevidal on Wed, 07 Jun 2006 15:38:09 GMT
View Forum Message <> Reply to Message

kir wrote on Wed, 07 June 2006 11:27Quote:* How fast (your best guess) is it over a gigabit network with a typical guest distro? Xen's live migration downtime is microseconds; when I read that OpenVZ has "zero downtime" migration, does that mean it's microseconds, too?

This is not really zero downtime -- it is zero downtime from the user's point of view, since user does not see it as a downtime. I guess the delay in processing is a few seconds (non-shared storage scenario). We are improving this time.

That's not too bad; so you're saying you depend upon the application (Samba, Apache, MySQL, PostgreSQL, NFS, etc.) to handle the brief hang without killing the connection? I know Samba and NFS can deal with this and I suppose the rest are OK...

I'm almost certain to choose OpenVZ, it looks like a great project.

Subject: Re: Xen vs. OpenVZ: live migration
Posted by kir on Wed, 07 Jun 2006 15:41:27 GMT
View Forum Message <> Reply to Message

Quote:so you're saying you depend upon the application (Samba, Apache, MySQL, PostgreSQL, NFS, etc.) to handle the brief hang without killing the connection? I know Samba and NFS can deal with this and I suppose the rest are OK...

This is actually a TCP feature. Default (and recommended by an RFC) TCP timeout is about two minutes if I remember correctly. So if the application is using TCP, it will be fine. Surely we tested that with Apache, MySQL, Oracle and other apps.

Subject: Re: Xen vs. OpenVZ: live migration
Posted by cdevidal on Wed, 07 Jun 2006 15:44:51 GMT
View Forum Message <> Reply to Message

kir wrote on Wed, 07 June 2006 11:41This is actually a TCP feature. Default (and recommended by an RFC) TCP timeout is about two minutes if I remember correctly. So if the application is using TCP, it will be fine. Surely we tested that with Apache, MySQL, Oracle and other apps.

ROCK ON!! I didn't know that...

And I suppose UDP doesn't care if some packets are lost?

Subject: Re: Xen vs. OpenVZ: live migration

Posted by kir on Wed, 07 Jun 2006 15:56:32 GMT
View Forum Message <> Reply to Message

Quote:And I suppose UDP doesn't care if some packets are lost?
That depends on the application which uses UDP. But UDP is unreliable protocol by default so apps should cope with the fact that some packets are lost.