Subject: Performance Question Posted by dietmar on Sat, 10 Mar 2007 19:53:35 GMT

View Forum Message <> Reply to Message

Hi all,

we have done some performance tests running our mail-gateway software on OpenVZ. OpenVZ is extremely fast compared to other solutions, but we still get a 20% performance loss compared to running native. Our benchmark put the server under extremely high load (50 postfix smtpd, 40 spamassassin, 40 postgres database connections, ...). The Benchmark measures the maximum throughput of mails (Mails/second).

Is that 20% loss normal, or is it possible to improve that?

- Dietmar

Subject: Re: Performance Question

Posted by rickb on Sun, 11 Mar 2007 01:18:44 GMT

View Forum Message <> Reply to Message

20% performance loss using Openvz isn't normal. Think of openvz as resource containers with tools to view container stats on the HN, and when a PID is in the context, it has many restrictions to "stay" in its context.

Other, true and heavier virtualization architectures such as vmware, xen, etc introduce an interface for the virtualized kernel to talk to the underlying HN. This is where a performance loss is induced.

20% is extremely high, and very unexpected.

Rick

Subject: Re: Performance Question

Posted by JimL on Sun, 11 Mar 2007 02:25:51 GMT

View Forum Message <> Reply to Message

This is totally off the wall, but is it possible your application is sensitive to the amount of memory available and you have allocated less for the VE?

Jim.

Subject: Re: Performance Question

Posted by Vasily Tarasov on Sun, 11 Mar 2007 12:38:35 GMT

View Forum Message <> Reply to Message

Hello,

we need some information from you:

- 1) What to what have you compared? I mean, have you compared VE0 to VE or VE to some other kernel (non OpenVZ)?
- 2) You should use absolutely the same binaries in VE and in VE0!
- 3) Disk layout is very pertinent.
- 4) Try to set all UBC limits to unlimited value. Even if you don't have failcounters it may be pertinent.

Vasily.

Subject: Re: Performance Question

Posted by dev on Sun, 11 Mar 2007 12:42:58 GMT

View Forum Message <> Reply to Message

What was the test?

- what environment do you refer to when refering to native? host system? Which OS/kernel etc.?
- during the test you used most likely different directories for storing emails etc. The fact is that HDD speed varies very much for disk beginning and end.

So to make results comparable you should do the following:

- a) create a separate partition of size which will be used totally up (plus some small reserve)
- b) mount this partition to /var/spool when testing host environment and /vz/root/VEID/var/spool when testing VE.

This will guarantee that the same physical sectors of HDD will be used in both tests.

- set VE UBC setting to unlimited. This will prevent resource shortages during the test.

There can be other tunings done and checked for (e.g. mounting /vz with noatime), but w/o test case and test procedure it is hard to advise.

Subject: Re: Performance Question

Posted by dietmar on Mon, 12 Mar 2007 12:45:44 GMT

View Forum Message <> Reply to Message

dev wrote on Sun, 11 March 2007 07:42What was the test?

You can find our OpenVZ template at:

http://www.proxmox.com/cms_proxmox/cms/upload/bittorrent/debian-3.1-proxmox-mailgateway.tar.gz

Its an OpenVZ template for our Mail Gateway software appliance.

The test was simply sending several hundred testmails to the system, doing virus an spam filtering. I can send you the test scripts.

Quote:

- what environment do you refer to when refering to native? host system? Which OS/kernel etc.?

debian/2.6.18-3

Quote:

- during the test you used most likely different directories for storing emails etc. The fact is that HDD speed varies very much for disk beginning and end.

So to make results comparable you should do the following:

- a) create a separate partition of size which will be used totally up (plus some small reserve)
- b) mount this partition to /var/spool when testing host environment and /vz/root/VEID/var/spool when testing VE.

This will guarantee that the same physical sectors of HDD will be used in both tests.

Do you think that explains 15-20% Performance loss? Anyways, I cant change that for several reasons.

Quote:

- set VE UBC setting to unlimited. This will prevent resource shortages during the test.

There can be other tunings done and checked for (e.g. mounting /vz with noatime), but w/o test case and test procedure it is hard to advise.

I already set UBC settings to unlimited, but still get 15% less performance (which is quite good compared to other VM technologies)

I wonder if someone else did performance tests on comparable systems (mail server, mail gateways) running on OpenVZ?

Subject: Re: Performance Question

Posted by dietmar on Mon, 12 Mar 2007 12:53:28 GMT

View Forum Message <> Reply to Message

Vasily Tarasov wrote on Sun, 11 March 2007 07:384) Try to set all UBC limits to unlimited value. Even if you don't have failcounters it may be pertinent.

I generated a reasonable default using vzsplit -n 1, then I set memory to 2GB:

VMGUARPAGES="524288:2147483647"

OOMGUARPAGES="524288:2147483647" PRIVVMPAGES="524288:2147483647"

or whats the best way to set no(unlimited) UBC limits?

- Dietmar

Subject: Re: Performance Question

Posted by dev on Mon, 12 Mar 2007 13:13:18 GMT

View Forum Message <> Reply to Message

dietmar wrote on Mon, 12 March 2007 15:45dev wrote on Sun, 11 March 2007 07:42What was the test?

You can find our OpenVZ template at:

http://www.proxmox.com/cms_proxmox/cms/upload/bittorrent/debian-3.1-proxmox-mailgateway.tar.gz

Its an OpenVZ template for our Mail Gateway software appliance.

The test was simply sending several hundred testmails to the system, doing virus an spam filtering. I can send you the test scripts.

It would be great, since then we could try to reproduce.

dietmar wrote on Mon, 12 March 2007 15:45

Quote:

- what environment do you refer to when refering to native? host system? Which OS/kernel etc.?

debian/2.6.18-3

Quote:

- during the test you used most likely different directories for storing emails etc. The fact is that HDD speed varies very much for disk beginning and end.

So to make results comparable you should do the following:

- a) create a separate partition of size which will be used totally up (plus some small reserve)
- b) mount this partition to /var/spool when testing host environment and /vz/root/VEID/var/spool when testing VE.

This will guarantee that the same physical sectors of HDD will be used in both tests.

Do you think that explains 15-20% Performance loss? Anyways, I cant change that for several reasons.

It can account easily for almost double performance variations in some cases. Though I don't

claim it is 100% the issue in your test case, just trying to provide more likely reasons for performance differences and which we can check for.

dietmar wrote on Mon, 12 March 2007 15:45

Quote:

- set VE UBC setting to unlimited. This will prevent resource shortages during the test.

There can be other tunings done and checked for (e.g. mounting /vz with noatime), but w/o test case and test procedure it is hard to advise.

I already set UBC settings to unlimited, but still get 15% less performance (which is quite good compared to other VM technologies) dietmar wrote on Mon, 12 March 2007 15:45

I'm pretty sure it can be optimized further out. It's just not that easy to find out the bottleneck remotely maybe it's possible to get an access to your testbed?

Do you use 100Mb or 1Gb link?

I wonder if someone else did performance tests on comparable systems (mail server, mail gateways) running on OpenVZ?

Subject: Re: Performance Question

Posted by dietmar on Mon, 12 Mar 2007 14:37:12 GMT

View Forum Message <> Reply to Message

Here ist the test script (perl). I tried to add the necessary documentation, but just ask if you have questions about howto setup the mail gateway.

- Dietmar

File Attachments

1) perftest.pl, downloaded 395 times

Subject: Re: Performance Question

Posted by dietmar on Mon, 12 Mar 2007 14:46:46 GMT

View Forum Message <> Reply to Message

If you test on native hardware you need to download and install our iso image (warning: formats the Harddisk when installing):

http://www.proxmox.com/cms_proxmox/cms/upload/bittorrent/pro xmox-mailgateway_1.7.iso (or the vmware appliance if you want to test that).

Also apply the hotfix after installation (use the web interface to apply the hotfix (Server/Updates). http://www.proxmox.com/cms_proxmox/cms/upload/bittorrent/hot fix_1.7-1702.bin

Subject: Re: Performance Question Posted by dev on Mon, 12 Mar 2007 16:10:49 GMT

View Forum Message <> Reply to Message

ok. will try to reproduce.