
Subject: How does one size a HN? Is RAM or CPU more important?

Posted by [jarcher](#) on Thu, 04 Jan 2007 06:49:40 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello...

I have some questions about sizing a hardware node. I read the OpenVZ docs and noted a few things.

First, it says that the low memory, which is 3.6GB when using the enterprise kernel or 382MB when using the Uniprocessor or SMP kernel, is the "most important" RAM area (page 59). Also, on page 60, it says that the number of VPSs can be doubled by using disk swap space.

I don't quite understand the implications of this. Does this mean that installing physical RAM beyond 4GB in the HN is a waste of money, or is that memory used by VPSs? Is this a matter of memory pages being swapped into and out of low memory as needed, and the RAM beyond the first 5.6GB is essentially non-disk swap?

Regarding CPU power, I see that OpenVZ has an app to calculate the "power" of the node in CPU units and we can control the CPU units a VPS is allocated. Is there anywhere a comparison of how many CPU units equates to a specific processor, even roughly? If I had a dual-xenon, how would I know how many VPSs I can cut this into?

I also saw in the manual that it is okay to over commit the HN, but I'm wondering how much over commitment we can get away with. I realize it depends upon how busy each VPS is and there is no hard and fast number for this, but I would appreciate any experience you folks could offer.

Finally, I saw in several places that a HN could be able to run 20 to 50 VPSs, but I'm wondering how much hardware needs to be put in the HN to accomplish that. What tends to be the limiting factor, RAM or CPU? I also seem to recall being told that 100 VPSs are possible on a single HN.

Thanks very much. I would appreciate any insight, as I will be ordering hardware soon.

Subject: Re: How does one size a HN? Is RAM or CPU more important?

Posted by [rickb](#) on Fri, 05 Jan 2007 00:41:33 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quote:physical RAM beyond 4GB in the HN is a waste of money

Ram above 4GB is not a waste. I have a HN with 16GB of memory and the VEs can use all of it.

Quote:

Regarding CPU power, I see that OpenVZ has an app to calculate the "power" of the node in CPU units and we can control the CPU units a VPS is allocated. Is there anywhere a comparison of how many CPU units equates to a specific processor, even roughly? If I had a dual-xenon, how

would I know how many VPSs I can cut this into?

Read my post here about cpuunits. Its an arbitrary value, it doesn't correlate to "power" .:

http://forum.openvz.org/index.php?t=tree&goto=8620&&srch=cpuunits+cpulimit#msg_8620

Quote: I'm wondering how much over commitment we can get away with.

Its totally dependent on your real usage needs. You can spawn thousands of VEs that do nothing and your system will never know the difference. Sleeping processes use very little resources. Or, you can have one running process per cpu in your server and it will choke. So, it totally depends on your applications. You can choke out any server with 1 heavy VE or load up 10,000 VEs which do almost nothing on a Pentium 3.

Quote: I would appreciate any insight

I have used all sorts of combinations from single cpu, dual, quad, xeon, opteron, dual cores on each, etc. Best bang for the buck in my opinion and from my experience

Late model supermicro board, 2x dual core 2.xGHz opteron, 8GB ECC, 4X disk (whatever you can afford, scsi better of course) @ Raid 1+0, hardware raid with 128/256MB cache enabled.

Rick Blundell

Subject: Re: How does one size a HN? Is RAM or CPU more important?

Posted by [rickb](#) on Sun, 07 Jan 2007 23:07:31 GMT

[View Forum Message](#) <> [Reply to Message](#)

If you are using SATA disks, I have found this to be a very good card:

3ware 9500S-4LP

It supports Raid 10, has a good linux CLI management tool (tw_cli), 256MB caching, seems to offload a nice amount of i/o processing to the card (compared to software raid for example), and is supported by all of the openvz kernels.

Rick Blundell

Subject: Re: How does one size a HN? Is RAM or CPU more important?

Posted by [alticon-brian](#) on Wed, 21 Feb 2007 06:04:35 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hey rick,

I just wanted to drop a optimization tip your way.

I've deployed about 10 database servers for clients now using an array of 9500S, 9550SX, & the new 9650 cards. After a lot of testing and a phone call to confirm my findings, the 3ware cards

actually perform better on RAID5/6.

The reason is, the XOR chip is optimized for doing parity on raid5. On the note of intel/amd was it a big difference? All we've deployed is the 5000 Series Xeon chips and they've done a great job.

--brian

Subject: Re: How does one size a HN? Is RAM or CPU more important?

Posted by [rickb](#) on Wed, 21 Feb 2007 06:19:38 GMT

[View Forum Message](#) <> [Reply to Message](#)

Brian, you mean to say that raid5/6 have better performance then raid10 on this card? What type of testing/environment are you running these comparisons in? The write speed for raid5 will never compare to raid10 to my knowledge, just by the design, without looking at the performance of the card/chip.

Looking forward to more info from you, because raid10 incurs a lot of disk space overhead of course, which I assumed was justified by the performance benefit of raid10.

Rick

Subject: Re: How does one size a HN? Is RAM or CPU more important?

Posted by [alticon-brian](#) on Wed, 21 Feb 2007 06:50:45 GMT

[View Forum Message](#) <> [Reply to Message](#)

Yup. It blew my mind. I started out with our servers running RAID 10 and things were o.k., moved to raid 5 (and am building new pxeboot images to support the raid6 card) and wowsers.

Here's a good whitepaper from 3ware on the topic:

http://www.3ware.com/kb/attachments/Linux2.6WP_0701.pdf

And a knowledgebase article on tuning from them:

<http://www.3ware.com/kb/article.aspx?id=11050>

these are the results:

All of our database servers are

supermicro 6025T

dual Xeon 5000 3.0Ghz Dual Core

8GB RAM

8 320GB Seagate Perpendicular Drives

3ware 9550SX/9650

Subject: Re: How does one size a HN? Is RAM or CPU more important?

Posted by [alticon-brian](#) on Wed, 21 Feb 2007 06:53:07 GMT

[View Forum Message](#) <> [Reply to Message](#)

oh, also...

using MD in linux RAID 10 would be faster.

the benefit comes from the optimized XOR chip on that card.

it's that chip that drives the price of those cards up to around \$500 for the 8 drive models.

basically, it comes down to software raid vs hardware raid. in this case, you're paying for the benefit of not only offloading the cycles to that card, but getting the space back from your extra disks in the array.

Subject: Re: How does one size a HN? Is RAM or CPU more important?

Posted by [rickb](#) on Wed, 21 Feb 2007 11:27:40 GMT

[View Forum Message](#) <> [Reply to Message](#)

Thank you for the information, truly a valuable post. I have some further questions if you have time.

What sector readahead value are you using with blockdev? Default on my card is 256 (blockdev --getra /dev/sdX), the pdf suggest 16384 secs.

I am using ext3 but am not using any tuned values beyond noatime. Have you found and performance gains from optimizations on the fs level?

I am using the cfq io scheduler, because after reading and unofficially testing it, it seems to deliver the best io response time for a shared resource system. Although deadline may offer better overall io performance, fairness is more relevant in a shared resource environment. What are your thoughts on this?

My servers with this card have 8GB of ram, and I have created 8GB of swap; after 63 days of uptime, the system stays constant with ~2.5GB of swap utilized. I have read that adjusting the amount of swap in the system can positively effect performance. Do you have any ideas there?

```
[root@gallium ~]# free -m
      total    used    free   shared  buffers   cached
Mem:      8029    8004     24      0       97    1586
-/+ buffers/cache:    6320    1709
Swap:      8000    2609   5391
```

My primary indicator of "disk performance" in a shared resource environment is a (lower) iowait

value over time (vmstat, sar, iostat, etc). Do you find this relevant?

-Rick Blundell

Subject: Re: How does one size a HN? Is RAM or CPU more important?

Posted by [jarcher](#) on Thu, 22 Feb 2007 06:31:26 GMT

[View Forum Message](#) <> [Reply to Message](#)

rickb wrote on Thu, 04 January 2007 19:41 I have used all sorts of combinations from single cpu, dual, quad, xeon, opteron, dual cores on each, etc. Best bang for the buck in my opinion and from my experience

Late model supermicro board, 2x dual core 2.xGHz opteron, 8GB ECC, 4X disk (whatever you can afford, scsi better of course) @ Raid 1+0, hardware raid with 128/256MB cache enabled.

Hi Rick, thanks very much for the response! I apologize for not replying earlier but I only now saw your reply.

Well I guessed well, because I just built up a SuperMicro dual socket Xeon with a pair of Xeon 5140 chips and 8GB RAM. Well, I am still waiting for the RAM to arrive but that will be Friday. I can add RAM up to 64GB and I can upgrade the CPUs if I need to to the quad-core. They will be cheaper by the time I need them.

I'll be using a SAN box via iSCSI for the storage.

The SAN box, BTW, has an Areca ARC-1230 RAID card. I have not measured it's performance, but from a feature and management standpoint I just love it. I'm running a 3TB array in RAID 6 for now and I have 6 empty sleds for expansion.

I'm hoping to be able to run at least 150 light to medium ('A' and 'B' machines from the example configurations) VPSes on this hardware. Fingers crossed.

Thanks!!

Subject: Re: How does one size a HN? Is RAM or CPU more important?

Posted by [jarcher](#) on Thu, 22 Feb 2007 06:54:36 GMT

[View Forum Message](#) <> [Reply to Message](#)

rickb wrote on Thu, 04 January 2007 19:41 Read my post here about cpuunits. Its an arbitrary value, it doesn't correlate to "power" .:

http://forum.openvz.org/index.php?t=tree&goto=8620&& amp; amp;srch=cpuunits+cpulimit#msg_8620

Does a dual-core processor also have hyperthreading in each core, or does a dual-core processor always count as 2 processors?

Subject: Re: How does one size a HN? Is RAM or CPU more important?

Posted by [rickb](#) on Thu, 22 Feb 2007 09:05:02 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quote: Does a dual-core processor also have hyperthreading in each core, or does a dual-core processor always count as 2 processors?

It depends on the model. Some do and some don't- check your intel specs sheet. I know this xeon is dual core with ht:

"Dual Core Processor 2.8GHz, 800MHz FSB, 2 x 2MB L2 cache, 135W, 90nm, mPGA604"

Quote: processor : 7
vendor_id : GenuineIntel
cpu family : 15
model : 4
model name : Genuine Intel(R) CPU 2.80GHz
stepping : 8
cpu MHz : 2794.252
cache size : 2048 KB
physical id : 1
siblings : 4
core id : 3
cpu cores : 2
fdiv_bug : no
hlt_bug : no
f00f_bug : no
coma_bug : no
fpu : yes
fpu_exception : yes
cpuid level : 5
wp : yes
flags : fpu vme de pse tsc msr pae mce cx8 apic mtrr pge mca cmov pat pse36 clflush dts
acpi mmx fxsr sse sse2 ss ht tm pbe lm pni monitor ds_cpl est cid xtp
bogomips : 5585.25

With this model, linux sees 8 CPUs. but there are two physical ones in there. Quad core, quad cpu, hyperthread. not sure about that one but sounds pretty crazy. 4*4*2

Subject: Re: How does one size a HN? Is RAM or CPU more important?
Posted by [jarcher](#) on Thu, 22 Feb 2007 14:50:32 GMT

[View Forum Message](#) <> [Reply to Message](#)

[quote title=rickb wrote on Thu, 22 February 2007 04:05]Quote:With this model, linux sees 8 CPUs. but there are two physical ones in there. Quad core, quad cpu, hyperthread.

Thanks again Rick, but I am still missing something. I recognize that output from cpuid, but what in there tells me it's a quad core? I see it said "cores 2" but what about HT?

Subject: Re: How does one size a HN? Is RAM or CPU more important?
Posted by [alticon-brian](#) on Thu, 22 Feb 2007 15:26:45 GMT

[View Forum Message](#) <> [Reply to Message](#)

FYI, Hyperthreading isn't really geared as a server technology. It's more for the desktop space.

This is especially the case where you have applications which actually benefit from multi-threading (i.e. databases) which perform much better with hyperthreading turned off.

As a rule, i turn off HT on all of our ovz boxes, but thats just because of previous experience and the knowledge that it shouldn't be getting some huge boost from it, not from any actual performance testing.

I'd love to hear the opinions of some other folks on this.
