Subject: poor performance in openvz kernels...
Posted by devonblzx on Fri, 13 Apr 2007 06:20:41 GMT
View Forum Message <> Reply to Message

I have a new Dual Quad-Core Clovertown (E5310) with a RAID5 and when I run unixbench (the wht variant) I get these results...

Default CentOS4 Kernel (2.6.9smp):
Start Benchmark Run: Thu Apr 12 22:01:07 EDT 2007
 22:01:07 up 1 min, 1 user, load average: 0.20, 0.08, 0.03

End Benchmark Run: Thu Apr 12 22:11:14 EDT 2007
 22:11:14 up 11 min, 1 user, load average: 27.25, 10.54, 4.39

TEST	BASELINE	RESULT	INDEX
	INDEX VALUES		
Dhrystone 2 using register variables	376783.7	73062884.2	1939.1
Double-Precision Whetstone	83.1	1452.6	174.8
Execl Throughput	188.3	11972.2	635.8
File Copy 1024 bufsize 2000 maxblocks	2672.0	90568.0	339.0
File Copy 256 bufsize 500 maxblocks	1077.0	22404.0	208.0
File Read 4096 bufsize 8000 maxblocks	15382.0	1112686.0	723.4
Pipe-based Context Switching	15448.6	1536903.0	994.8
Pipe Throughput	111814.6	7310880.6	653.8
Process Creation	569.3	48057.5	844.2
Shell Scripts (8 concurrent)	44.8	2787.1	622.1
System Call Overhead	114433.5	8538136.8	746.1
		=========	
FINAL SCORE			586.1

Kernel 2.6.18-stab027-smp:
Start Benchmark Run: Thu Apr 12 21:29:33 EDT 2007
 21:29:33 up 1 min, 1 user, load average: 0.22, 0.11, 0.04

End Benchmark Run: Thu Apr 12 21:39:43 EDT 2007
 21:39:43 up 11 min, 1 user, load average: 25.45, 10.04, 4.23

TEST	BASELINE	RESULT	INDEX
	INDEX VALUES		
Dhrystone 2 using register variables	376783.7	71079686.4	1886.5
Double-Precision Whetstone	83.1	1456.2	175.2
Execl Throughput	188.3	4733.3	251.4
File Copy 1024 bufsize 2000 maxblocks	2672.0	120019.0	449.2
File Copy 256 bufsize 500 maxblocks	1077.0	32769.0	304.3

File Read 4096 bufsize 8000 maxblocks	15382.0	1125105.0	731.4
Pipe-based Context Switching	15448.6	415588.7	269.0
Pipe Throughput	111814.6	6285422.1	562.1
Process Creation	569.3	11449.0	201.1
Shell Scripts (8 concurrent)	44.8	1505.2	336.0
System Call Overhead	114433.5	8894040.1	777.2
	=========		
FINAL SCORE		416.8	

Any idea why it would be such a big performance loss? It seems like the system has trouble with Pipe-based Context Switching, Process Creation, Shell Scripts and Execl Throughput.

I know that 2.6.18 has a problem with the CPU scheduler but I tested it on the rhel4 stable openvz 2.6.9 and turned up even worse results (around 350 final score).

Subject: Re: poor performance in openvz kernels...
Posted by dev on Fri, 13 Apr 2007 16:25:22 GMT
View Forum Message <> Reply to Message

have you used binary 2.6.18 kernel from openvz.org or some other kernel?
(openvz.org doesn't have "-smp" kernel)

Your results looks strange, we did comparison of UnixBench on one of our machines only a week ago and it wasn't that awful
So I guess it maybe related to:
- quad-core? unlikely, but still...
- your .config file?
- something else which should be investigated

We will rerun it on Monday on our side as well.

Subject: Re: poor performance in openvz kernels...
Posted by devonblzx on Fri, 13 Apr 2007 16:42:02 GMT
View Forum Message <> Reply to Message

I have used every different kind of kernel with similar results, I have used customized versions of 2.6.9 and 2.6.18 along with the binary versions with similar results. They all perform with a difference of 200 compared to the non-openvz kernels.

On a custom 2.6.18 (non-ovz) I received over 600, but with any openVZ 2.6.18 I cannot get over 420.

Subject: Re: poor performance in openvz kernels...
Posted by dev on Fri, 13 Apr 2007 16:47:28 GMT
View Forum Message <> Reply to Message

Will it be possible to get an access to your machine for investigating this?

Subject: Re: poor performance in openvz kernels...
Posted by devonblzx on Fri, 13 Apr 2007 16:50:53 GMT
View Forum Message <> Reply to Message

Would you be able to do it today? I'm actually planning on shipping this off for production, it is at my home right now.

Subject: Re: poor performance in openvz kernels...
Posted by devonblzx on Fri, 13 Apr 2007 16:56:19 GMT
View Forum Message <> Reply to Message

I have sent you the login information to dev@openvz.org.

I believe that is the email address you told me last time, right?

Subject: Re: poor performance in openvz kernels...
Posted by dev on Fri, 13 Apr 2007 17:10:00 GMT
View Forum Message <> Reply to Message

Yes, it's mine address. Will try today , though it's not that much time left :/

Subject: Re: poor performance in openvz kernels...
Posted by devonblzx on Fri, 13 Apr 2007 17:11:36 GMT
View Forum Message <> Reply to Message

Yes the time difference from Moscow to the US is sort of a pain

Thanks for trying atleast, feel free to run any tests you need...all 8 cores are yours for testing

Sincerely,
Devon

Subject: Re: poor performance in openvz kernels...
Posted by Alexandr Andreev on Fri, 20 Apr 2007 13:23:25 GMT
View Forum Message <> Reply to Message

Hello devonblzx:

1. Can you attach your:

1.1. /proc/cpuinfo
1.2. 'ps -ax' output before you start the unixbench test on CENTos kernel and on OVZ kernel also. Are you sure there are no other running processes while unixbench is running? I think it's better to run "/sbin/init 1" before start tests.

2. Can you also set "maxcpus=2" kernel command line option for OVZ kernel and rerun tests again? Then you can try "maxcpus=1".

We need some information to investigate this issue. It can be SMP-scalability problem and it can depends on your CPU's hardware topology.

Subject: Re: poor performance in openvz kernels...
Posted by devonblzx on Fri, 20 Apr 2007 16:24:36 GMT
View Forum Message <> Reply to Message

I think it is something with the eight cores...my server is in the process of being setup for production.

My cpuinfo just shows eight E5310 cores...

But I did receive about the same scores between the OVZ and CentOS kernels when I ran the regular 4-threaded test, but when I tried running the 8-threaded test is when I saw such a big difference, so I'm sure it is something with the scalability for more than 4-cores.

Subject: Re: poor performance in openvz kernels...
Posted by dev on Wed, 02 May 2007 10:51:00 GMT
View Forum Message <> Reply to Message

BTW, 2.6.18+ kernels have still CONFIG_UBC_DEBUG_KMEM=y which affects this microbenchmark pretty badly.
So if you have a chance, please recheck without it.

on 2.6.18-028stab031.1 kernel and 16 CPUs Intel machine we get the following:

kernel	030.1	RHEL5	030.1 (*)
FINAL SCORE	450.4	471.7	481.7

(*) - it is 030.1 kernel with CONFIG_UBC_DEBUG_KMEM=n

Subject: Re: poor performance in openvz kernels...
Posted by devonblzx on Thu, 10 May 2007 23:55:36 GMT
View Forum Message <> Reply to Message

Wow,

I didn't get to test this sooner, but as soon as I disabled the debug in the kernel the benchmark went from 380 to 660. So a big difference on it.

Looks like this one is solved.

Thanks for your help.

Subject: Re: poor performance in openvz kernels...
Posted by devonblzx on Thu, 24 May 2007 18:45:14 GMT
View Forum Message <> Reply to Message

Nevermind, I actually forgot to patch the kernel...

I ran it without the OVZ patch and the 2.6.18 kernel received 660...as soon as I patched it with OpenVZ it returned a score of 410.

Still not sure what the problem is.

Subject: Re: poor performance in openvz kernels...
Posted by dev on Fri, 25 May 2007 08:07:36 GMT
View Forum Message <> Reply to Message

well, we are no sitting w/o doing anything and the recent kernels in -git have some optimizations which should help a bit.
At least on 16 CPU we were able to gain a good results (+/- 1-2% of original scores). This kernel will be released soon.

Subject: Re: poor performance in openvz kernels...
Posted by hzqbbc on Sun, 01 Nov 2009 14:34:33 GMT
View Forum Message <> Reply to Message

This post is very old (May 2007) but we still encounter the performance problem.

Unixbench-4.1.0-wht-2 result different from 2.6.18-128.2.1.el5.028stab064.7 and 2.6.18-128.el5 (standard RHEL5 kernel), the difference is huge.

After searching openvz forum, i think i'm closing to the answer, but not sure, the huge perf loss may due to openvz UBC accounting.

So my question is *WHEN* openvz dev team will fix this problem? dev said he would release new kernel to archive better perf but i try the latest 2.6.18-128.2.1.el5 (stab064.7) and got no improvement.

My bench result:

```
========================================================================
BYTE UNIX Benchmarks (Version 4.1-wht.2, 8 threads)
System -- Linux vz 2.6.18-128.el5 #1 SMP Wed Jan 21 10:41:14 EST 2009 x86_64 x86_64
x86_64 GNU/Linux
                  65672880   3609080  58673952   6% /


End Benchmark Run: Mon Nov  2 06:18:21 CST 2009
 06:18:21 up 13 min,  2 users,  load average: 31.71, 12.44, 5.31


                    INDEX VALUES
TEST                                  BASELINE     RESULT      INDEX

Dhrystone 2 using register variables        376783.7 36719630.6      974.6
Double-Precision Whetstone                     83.1     1298.6      156.3
Execl Throughput                              188.3    20191.8     1072.3
File Copy 1024 bufsize 2000 maxblocks        2672.0   177441.0      664.1
File Copy 256 bufsize 500 maxblocks          1077.0    43517.0      404.1
File Read 4096 bufsize 8000 maxblocks       15382.0  1771809.0     1151.9
Pipe-based Context Switching                15448.6  2739102.5     1773.0
Pipe Throughput                            111814.6 11210018.3     1002.6
Process Creation                              569.3    64281.4     1129.1
System Call Overhead                       114433.5 19818827.5     1731.9
                                                                 =========
    FINAL SCORE                                                      840.3


========================================================================
BYTE UNIX Benchmarks (Version 4.1-wht.2, 8 threads)
System -- Linux vz 2.6.18-128.2.1.el5.028stab064.7 #1 SMP Wed Aug 26 15:47:17 MSD 2009
x86_64 x86_64 x86_64 GNU/Linux
                  65672880   3648024  58635008   6% /
```

Start Benchmark Run: Mon Nov 2 06:22:56 CST 2009
 06:22:56 up 1 min, 1 user, load average: 2.20, 0.88, 0.32

End Benchmark Run: Mon Nov 2 06:33:02 CST 2009
 06:33:02 up 11 min, 2 users, load average: 27.90, 11.13, 4.73

TEST	BASELINE	RESULT	INDEX
	INDEX VALUES		
Dhrystone 2 using register variables	376783.7	36500265.7	968.7
Double-Precision Whetstone	83.1	1295.2	155.9
Execl Throughput	188.3	3618.3	192.2
File Copy 1024 bufsize 2000 maxblocks	2672.0	167840.0	628.1
File Copy 256 bufsize 500 maxblocks	1077.0	41757.0	387.7
File Read 4096 bufsize 8000 maxblocks	15382.0	1752507.0	1139.3
Pipe-based Context Switching	15448.6	2243795.1	1452.4
Pipe Throughput	111814.6	14018287.3	1253.7
Process Creation	569.3	9315.9	163.6
Shell Scripts (8 concurrent)	44.8	1147.4	256.1
System Call Overhead	114433.5	20733248.3	1811.8
		=========	
FINAL SCORE			538.9

Subject: Re: poor performance in openvz kernels...
Posted by hzqbbc on Sat, 14 Nov 2009 14:31:32 GMT
View Forum Message <> Reply to Message

Anybody can explain why and provide a solution ?

The problem may exist in all openVZ kernel, hope offical development team can focus on it,
Thanks!
