Subject: *SOLVED* network performance
Posted by ep1p on Thu, 07 Jun 2007 18:55:25 GMT
View Forum Message <> Reply to Message

Hi,
I have three boxes running openvz on gentoo, and when trying to diagnose a problem with network performance to a seperate host not on my network, I found that the openvz boxes appear to get better performance than all other boxes I tested (to an order of perhaps 3x faster when wget'ing a file). This occurs both in the VE and on the host system.

The data file appears to be 50meg worth of A's so I guess is fairly compressable. What clever stuff is happening such that these boxes perform better? Kernel is 2.6.18-028stab031

Cheers,

john

Subject: Re: network performance
Posted by rickb on Fri, 08 Jun 2007 07:12:54 GMT
View Forum Message <> Reply to Message

This sounds rather fantastic !, but maybe too good.

Do you have the same performance (3x faster) if you reboot the same server into a non-openvz kernel?

I would say its more likely to be a minor hardware or physical layout difference between your servers (between the one that is 3x fast and 1x fast).

Subject: Re: network performance
Posted by dev on Fri, 08 Jun 2007 08:03:28 GMT
View Forum Message <> Reply to Message

I wish all users were wondering about such issues

Subject: Re: network performance
Posted by ep1p on Thu, 14 Jun 2007 00:34:56 GMT
View Forum Message <> Reply to Message

Hmm. I have several of these types of boxes so they should be a pretty much identical architecture (apart from maybe ag1 below). Some have single CPU, some have dual, some have a 3.2ghz chip instead of 2.8ghz and some have more RAM than others. The network cards are pretty much the same- intel onboard gigabit cards (connected to 100meg switches).

my openvz boxes:

Linux box1 2.6.18-028stab031 #1 SMP Thu May 10 11:26:56 GMT 2007 i686 Intel(R) Xeon(TM) CPU 3.20GHz GenuineIntel GNU/Linux

00:13:02 (4.15 MB/s)

Linux box2 2.6.18-028stab031 #1 SMP Fri May 11 14:27:07 GMT 2007 i686 Intel(R) Xeon(TM) CPU 3.20GHz GenuineIntel GNU/Linux

00:06:51 (4.48 MB/s)

Linux box3 2.6.20-gentoo-r8 #3 SMP Fri May 25 00:32:12 BST 2007 i686 Intel(R) Xeon(TM) CPU 3.20GHz GenuineIntel GNU/Linux

01:12:35 (5.22 MB/s)

the others:

Linux ag1 2.6.18-gentoo-r3 #1 SMP Thu Nov 30 06:46:32 GMT 2006 i686 Intel(R) Xeon(R) CPU 5160 @ 3.00GHz GenuineIntel GNU/Linux

01:12:02 (5.21 MB/s)

Linux ag2 2.6.14-hardened-r3 #4 SMP Wed Feb 22 15:03:41 GMT 2006 i686 Intel(R) Xeon(TM) CPU 2.80GHz GenuineIntel GNU/Linux

01:08:45 (1.26 MB/s)

Linux ag3 2.6.14-hardened-r6 #2 SMP Tue Apr 11 16:49:45 GMT 2006 i686 Intel(R) Xeon(TM) CPU 2.80GHz GenuineIntel GNU/Linux

01:19:32 (1.25 MB/s)

Linux ag4 2.6.20-hardened-r2 #3 SMP Tue Jun 12 16:35:52 BST 2007 i686 Intel(R) Xeon(TM) CPU 2.80GHz GenuineIntel GNU/Linux

02:20:04 (1.70 MB/s)

Linux ag5 2.6.16-hardened-r6 #2 SMP Mon May 22 21:09:25 BST 2006 i686 Intel(R) Xeon(TM) CPU 3.20GHz GNU/Linux

01:23:02 (1.22 MB/s)

Linux ag6 2.6.16-hardened-r6 #2 SMP Mon May 22 21:09:44 BST 2006 i686 Intel(R) Xeon(TM) CPU 2.80GHz GNU/Linux

01:25:23 (1.26 MB/s)

box3 was running the openvz kernel and also had similar faster speeds with that.

I originally noticed this as I saw what appeared to be a 10Mbps bottleneck on a connection to something I don't control, but found from a vps server on my network there wasn't such an issue. Speeds to a local 100meg source are pretty much identical and are as would be expected on all these boxes.

The network equipment my side isn't a problem, nor is routing, or IP address as I've tried putting the same IP on different box with no difference to results.

The only thing I can think of is kernel, network driver (within kernel) or other network parameters like MTU/MSS/ window scaling etc that would adjust how things work over distance. Are there any performance tweaks done to net stuff for openvz?

Any thoughts of where to look? I guess this looks less like of an openvz issue now though

<edit - perhaps its hardened kernel? hmm..>

john

Subject: Re: network performance
Posted by dev on Thu, 14 Jun 2007 07:29:37 GMT
View Forum Message <> Reply to Message

Quote:
<edit - perhaps its hardened kernel? hmm..>

I think that's the best guess
not sure what hardened kernel is though...
Is it some special gentoo kernel?

Subject: Re: network performance
Posted by rickb on Thu, 14 Jun 2007 07:34:41 GMT
View Forum Message <> Reply to Message

have you booted into the non-vz kernel like I suggested? It could be something as silly as the ethernet wire, different interrupt setup, something very minor yet physically different. I really doubt openvz can boost the network speed by a factor of 3.

Subject: Re: network performance
Posted by dev on Thu, 14 Jun 2007 07:41:59 GMT
View Forum Message <> Reply to Message

yes, but hardened linux can easily decrease performance by a factor of 3.

Subject: Re: network performance
Posted by Alexandr Andreev on Thu, 14 Jun 2007 09:05:51 GMT
View Forum Message <> Reply to Message

It would be interesting to check and compare network performance with ttcp programm. You can use gentoo ttcp package, or compile ttcp by hands:

http://redvip.homelinux.net/varios/ttcp.c

Usage:

1. Test UDP performance (20000 60K packets):
receiver:
ttcp -r -s -u -l 60000
transmitter:
ttcp -t -s -u -l 60000 -n 20000 receiver.ip.host

2. TCP performance (20000 60K packets):
receiver:
ttcp -r -s -l 60000
transmitter:
ttcp -t -s -l 60000 -n 20000 receiver.ip.host

==========

In theory, ttcp must show approximately ~1Gb/sec (125 MB/s) for gigabit links. I hope on OpenVZ you will see 115-120 MB/s. But gentoo case is much more interesting, can you test gentoo kernel?

Subject: Re: network performance
Posted by adeeln on Sat, 07 Jul 2007 23:00:42 GMT
View Forum Message <> Reply to Message

I think the network speedup you're referring to might be because of the gentoo specific patches applied to the kernel. But overall, I've found Gentoo to be a faster OS than most others I had tried in the past (i've been using Gentoo for about 4 years now). Granted, I haven't tried any other distro recently, but then again, I don't feel that i have a real need to. Just my 2 cents

Subject: Re: network performance
Posted by ep1p on Sun, 08 Jul 2007 00:01:17 GMT
View Forum Message <> Reply to Message

It did turn out to be something in the gentoo hardened kernel that slows it down. Don't know what exactly, but removing that from both ends fixed the problem. I've still not worked out why it only causes a problem over a certain distance though. Perhaps gentoo hardened limits MTU or TCPMSS size or something.

john

Subject: Re: network performance
Posted by ep1p on Tue, 07 Aug 2007 21:36:19 GMT
View Forum Message <> Reply to Message

Infact, I have narrowed this down incase anyone is interested.

/proc/sys/net/core/rmem_max
/proc/sys/net/core/wmem_max
and
/proc/sys/net/ipv4/tcp_rmem
/proc/sys/net/ipv4/tcp_wmem

values were causing the issues. I.e. tcp send and receive buffers. The first two are the limit, the second two are the linux auto-tuning tcp buffers. You want the buffer size set to twice the bandwidth delay product. That means,

buffer size = bandwidth * round trip time (as reported by ping).

The gentoo hardened kernel I was using set the max to about 130k, where as the openvz kernel patch to the gentoo kernel set the autotuning max to 655360. For the link I'm using the above buffer size calc recons I need about 500k so hence the openvz kernel showed better spepeds.

To adjust the buffers, stick the variables in /etc/sysctl.conf like this:
net.ipv4.tcp_rmem = 4096 87380 2076672
net.ipv4.tcp_wmem = 4096 16384 2076672

then running 'sysctl -p' will apply them.

Hope this is of use to someone else.

Cheers,

john

Subject: Re: network performance
Posted by dev on Wed, 08 Aug 2007 07:43:48 GMT
View Forum Message <> Reply to Message

thanks you very much for the info!