
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
