
Subject: Bandwidth limiting crashes the machine

Posted by [eugenio pacheco](#) on Wed, 03 May 2006 20:43:32 GMT

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

Hi,

I'm running openvz on a Celeron 3066 with 1GB RAM and 160GB of disk space and I wished to limit traffic speed of a VPS. I searched the web and I came to a very interesting website:

<http://lartc.org/howto/>

They offer a script (the script can be found at

<http://lartc.org/howto/lartc.cookbook.ultimate-tc.html#AEN2233> and I also added it later on) that limits download and upload speed by using htb. The problem is that when I use this script on the host, after some time (a day or two) the server crashes. I don't know if it's only the network connection that stops, or the entire machine freezes (it's a remote server), but I do know that the computer just stops answering. Can someone help me?

```
-----  
#!/bin/bash
```

```
# The Ultimate Setup For Your Internet Connection At Home
```

```
#
```

```
#
```

```
# Set the following values to somewhat less than your actual download
```

```
# and uplink speed. In kilobits
```

```
DOWNLINK=1024
```

```
UPLINK=1024
```

```
DEV=eth0
```

```
# clean existing down- and uplink qdiscs, hide errors
```

```
tc qdisc del dev $DEV root 2> /dev/null > /dev/null
```

```
tc qdisc del dev $DEV ingress 2> /dev/null > /dev/null
```

```
##### uplink
```

```
# install root CBQ
```

```
tc qdisc add dev $DEV root handle 1: cbq avpkt 1000 bandwidth 10mbit
```

```
# shape everything at $UPLINK speed - this prevents huge queues in your
```

```
# DSL modem which destroy latency:
```

```
# main class
```

```
tc class add dev $DEV parent 1: classid 1:1 cbq rate ${UPLINK}kbit \
```

```
allot 1500 prio 5 bounded isolated
```

```
# high prio class 1:10:
```

```

tc class add dev $DEV parent 1:1 classid 1:10 cbq rate ${UPLINK}kbit \
  allot 1600 prio 1 avpkt 1000

# bulk and default class 1:20 - gets slightly less traffic,
# and a lower priority:

tc class add dev $DEV parent 1:1 classid 1:20 cbq rate [9*$UPLINK/10]kbit \
  allot 1600 prio 2 avpkt 1000

# both get Stochastic Fairness:
tc qdisc add dev $DEV parent 1:10 handle 10: sfq perturb 10
tc qdisc add dev $DEV parent 1:20 handle 20: sfq perturb 10

# start filters
# TOS Minimum Delay (ssh, NOT scp) in 1:10:
tc filter add dev $DEV parent 1:0 protocol ip prio 10 u32 \
  match ip tos 0x10 0xff flowid 1:10

# ICMP (ip protocol 1) in the interactive class 1:10 so we
# can do measurements & impress our friends:
tc filter add dev $DEV parent 1:0 protocol ip prio 11 u32 \
  match ip protocol 1 0xff flowid 1:10

# To speed up downloads while an upload is going on, put ACK packets in
# the interactive class:

tc filter add dev $DEV parent 1: protocol ip prio 12 u32 \
  match ip protocol 6 0xff \
  match u8 0x05 0x0f at 0 \
  match u16 0x0000 0xffc0 at 2 \
  match u8 0x10 0xff at 33 \
  flowid 1:10

# rest is 'non-interactive' ie 'bulk' and ends up in 1:20

tc filter add dev $DEV parent 1: protocol ip prio 13 u32 \
  match ip dst 0.0.0.0/0 flowid 1:20

##### downlink #####
# slow downloads down to somewhat less than the real speed to prevent
# queuing at our ISP. Tune to see how high you can set it.
# ISPs tend to have *huge* queues to make sure big downloads are fast
#
# attach ingress policer:

tc qdisc add dev $DEV handle ffff: ingress

# filter *everything* to it (0.0.0.0/0), drop everything that's

```

coming in too fast:

```
tc filter add dev $DEV parent ffff: protocol ip prio 50 u32 match ip src \
  0.0.0.0/0 police rate ${DOWNLINK}kbit burst 10k drop flowid :1
-----
```

Subject: Re: Bandwidth limiting crashes the machine
Posted by [Vasily Tarasov](#) on Thu, 04 May 2006 08:17:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

Can you, please, specify what openVZ kernel version do you use?
Thanks.

Subject: Re: Bandwidth limiting crashes the machine
Posted by [eugeniopacheco](#) on Thu, 04 May 2006 09:48:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

I'm very sorry about that, I completely forgot the kernel, here goes all information on the machine:

```
[root@t248 ~]# uname -a
Linux t248.1paket.com 2.6.8-022stab072.2 #1 Mon Mar 20 14:19:14 MSK 2006 i686 i686 i386
GNU/Linux
```

```
[root@t248 ~]# vzlist
  VPSID   NPROC STATUS IP_ADDR   HOSTNAME
  101     73 running x.x.x.x   x
  102     66 running x.x.x.x   x
  103     23 running x.x.x.x   x
```

```
[root@t248 ~]# vzmemcheck
Output values in %
LowMem LowMem  RAM MemSwap MemSwap Alloc Alloc Alloc
util commit util util commit util commit limit
  1.57 28.58 15.31 15.31 28.04 81.91 28.04 194.86
```

```
[root@t248 ~]# vzmemcheck -A
Output values in Mbytes
LowMem LowMem  RAM MemSwap MemSwap Alloc Alloc Alloc
util commit util util commit util commit limit
  5.12 93.45 153.22 153.22 280.68 819.98 280.68 1950.67
326.00 326.00 1001.00 1001.00 1001.00 1001.00 1001.00 1001.00
```

```
[root@t248 ~]# vzmemcheck -Av
Output values in Mbytes
veid   LowMem LowMem  RAM MemSwap MemSwap Alloc Alloc Alloc
      util commit util util commit util commit limit
  101   2.86 34.33 57.65 57.65 96.74 187.33 96.74 659.33
  102   1.18 24.79 58.94 58.94 87.20 495.98 87.20 845.11
```

103 1.09 34.33 36.64 36.64 96.74 136.68 96.74 446.24

Summary: 5.12 93.45 153.22 153.22 280.68 819.98 280.68 1950.67
326.00 326.00 1001.00 1001.00 1001.00 1001.00 1001.00 1001.00

```
[root@t248 ~]# vzcpucheck
Current CPU utilization: 42664
Power of the node: 151267
[root@t248 ~]# cat /proc/cpuinfo
processor      : 0
vendor_id    : GenuineIntel
cpu family   : 15
model        : 4
model name   : Intel(R) Celeron(R) CPU 3.06GHz
stepping     : 1
cpu MHz      : 3061.389
cache size   : 256 KB
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 sep mtrr pge mca cmov
pat pse36 clflush dts acpi mmx fxsr sse sse2 ss ht tm pbe pni monitor ds_cpl cid
bogomips     : 6045.69
```

Also, if someone could please let me know where the virtual memory comes from I would much appreciate, for when I run `vzmemcheck -Av` it shows limit numbers at the bottom, and it shows that the limit on my server for `alloc util` is 1001, same as RAM, but the RAM is way underused, so the question is, can I use more `alloc` memory than 1001 or it is limited as the RAM is?

Thanks for all your help.

Regards,

Eugenio Pacheco

Subject: Re: Bandwidth limiting crashes the machine
Posted by [Vasily Tarasov](#) on Fri, 05 May 2006 10:14:46 GMT
[View Forum Message](#) <> [Reply to Message](#)

Some comments.

In the script applied you use interface `eth0`, so you run it on the host machine. Then traffic control

affects all VPSs. Do you really wish such effect?

Inside VPS you can use interface `venet0`, to control traffic of this VPS only. But, of course, in this case, VPS root can change traffic control...

As concerns node crash. It's really necessary to know symptoms more precisely, `/var/log/messages`, etc. We even don't know was it crash or some network problem...

> Also, if someone could please let me know where the virtual memory comes from I would much appreciate, for when I run `vzmemcheck -Av` it shows limit numbers at the bottom, and it shows that the limit on my server for `alloc util` is 1001, same as RAM, but the RAM is way underused, so the question is, can I use more alloc memory than 1001 or it is limited as the RAM is?

`vzmemcheck` is rather strange utility on my mind =)

Bottom string is not limits actually, it's just a real `lowmem`, `totalmem` and `swap` values of your node (with some coeff.)

I suggest you to use `/proc/user_beancounters` file, there you can see limits, barriers, holds in real units. Also "OpenVZ User's Guide" contains information about this file.

Thanks.

Subject: Re: Bandwidth limiting crashes the machine
Posted by [eugenipacheco](#) on Sat, 06 May 2006 10:22:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

Thanks a lot for all the answers, it really helped me.

Right after I posted this problem with bandwidth limit, I also thought of using the device `venet0`, but it only limited the incoming bandwidth, not the outgoing, so I had the idea of making another script and it's running 3 days ago with no problems (at least up to now) Here it goes:

```
#!/bin/bash
```

```
DEV=eth0
```

```
tc qdisc del dev $DEV root
tc qdisc add dev $DEV root handle 1: cbq avpkt 1000 bandwidth 100mbit
tc class add dev $DEV parent 1: classid 1:1 cbq rate 1024kbit allot 1500 prio 5 bounded isolated
tc filter add dev $DEV parent 1: protocol ip prio 16 u32 match ip src x.x.x.x flowid 1:1
tc qdisc add dev $DEV parent 1:1 sfq perturb 10
```

```
DEV2=venet0
```

```
tc qdisc del dev $DEV2 root
tc qdisc add dev $DEV2 root handle 1: cbq avpkt 1000 bandwidth 100mbit
tc class add dev $DEV2 parent 1: classid 1:1 cbq rate 1024kbit allot 1500 prio 5 bounded isolated
tc filter add dev $DEV2 parent 1: protocol ip prio 16 u32 match ip dst x.x.x.x flowid 1:1
tc qdisc add dev $DEV2 parent 1:1 sfq perturb 10
```

This script limits the bandwidth for the ip x.x.x.x to 1024kbit/s both incoming and outgoing. It really works... Even if a client of yours use their vps to ddos, it will be stopped at 1024kbit/s, so if your machine has a 100mbit port it won't even affect your machine. Even if he gets ddos, it will also slow down the packages that comes to the machine and goes to the VPS. WARNING: From my own experience, I got one ip that was being ddosed, and it was slowing the entire VPS down, only the VPS, for it was limited. The problem is, I made the most stupid thing and tried to delete the ip address. It wasn't a good choice, for the incoming ddos now was going to the host machine, for the ip address was already routed to go through the host machine. The packets went to the host machine and stopped there since they couldn't find the ip address they were originally going to. RESULT: the entire machine was affected... So if you guys get ddosed to 1 ip, just limit the bandwidth to 32kbit/s let's say and ask the DC to block it on the router, DO NOT delete the ip address or your entire machine will be affected.

Now, there is another script I'm using to check bandwidth used (incoming and outgoing). It's by using ip tables and thanks to someone else that have posted it.

You run it on the host machine...

```
#!/bin/bash
DEV=eth0
iptables -A FORWARD -o $DEV -s x.x.x.x
iptables -A FORWARD -i $DEV -d x.x.x.x
```

This will set the iptables to log the bandwidth used. Then it can be seen by using:

```
#!/bin/bash
iptables -L FORWARD -v -x
iptables -L FORWARD -v
```

The first line shows the real numbers in bytes, expanded. The second one shows the numbers in Kbyte or Mbyte... It will show something like:

```
pkts-bytes-target-prot-opt-in-out---source---destination
117K-18M-----all- -- -any-eth0-x.x.x.x--anywhere
114K-17M-----all- -- -eth0-any-anywhere-x.x.x.x
```

As you can see the second line shows outgoing bandwidth while the third line shows incoming bandwidth.

If you want to reset the counters, just use:
iptables -Z

Hope this helped... And thanks for everything:)

Regards,

Eugenio Pacheco

Subject: Re: Bandwidth limiting crashes the machine
Posted by [jbond007](#) on Sat, 11 Aug 2007 11:11:14 GMT

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

dear frind this error make me crazy
[root@unknown ~]# tc qdisc del dev \$DEV root
RTNETLINK answers: No such file or directory
[root@unknown ~]#

please help me
Thank you

Subject: Re: Bandwidth limiting crashes the machine
Posted by [vaverin](#) on Sat, 11 Aug 2007 11:44:34 GMT

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

This command required to drop old traffic shapping settings. However if traffic control has not been configured before -- you get this answer.
Just ignore it and run other commands.

Subject: Re: Bandwidth limiting crashes the machine
Posted by [jbond007](#) on Sat, 11 Aug 2007 12:20:55 GMT

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

dear so if have RTNETLINK answers: No such file or directory
i must ignore?
Thank you for your help'
please reply!

how limit only 254 kb

Subject: Re: Bandwidth limiting crashes the machine
Posted by [vaverin](#) on Sat, 11 Aug 2007 12:38:38 GMT

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

jbond007 wrote on Sat, 11 August 2007 16:20dear so if have RTNETLINK answers: No such file or directory
i must ignore?

Yes, if "tc qdisc del dev \$DEV root" returns
"RTNETLINK answers: No such file or directory"
you can ignore it. This message means that kernel do not have old settings for \$DEV device, and
you can apply new settings.

Subject: Re: Bandwidth limiting crashes the machine
Posted by [jbond007](#) on Sat, 11 Aug 2007 14:45:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thank you for your help

how we test if work ?

how i see the changes that i been made ?
Thank you
