

Subject: High latency on first connection to VE
Posted by [avser](#) on Thu, 29 Sep 2011 20:12:56 GMT
[View Forum Message](#) <> [Reply to Message](#)

Help needed!

Issue with first packet latency. This issue happens both on TCP and UDP.

if a connection is attempted the first packet experiences huge latency causing application issues
immediate subsequent attempts do not have the problem

after some idle period the problem occurs again

isshe is very consistent and happens across all our servers in production and in dev envirohments not only the example machines below

unfortunately this firts attempt is very important for our setup, please, let us know your thoughts

SETUP:

host1 is a machine from which ping is made and IP 172.16.0.20 on eth0

host2 is a host node for VE 1000 with IP 172.16.0.8 on eth0

both eht0 interfaces are connected via gigabit L2 switch

VE 1000 is on host2 on a venet interface and IP 172.16.0.130

```
[root@host2 ~]# cat /etc/redhat-release
```

CentOS release 5.4 (Final)

```
[root@host2 /]# uname -a
```

```
Linux avs02 2.6.18-164.10.1.el5.028stab067.4 #1 SMP Thu Jan 14 21:23:12 MSK 2010 x86_64
```

x86_64 x86_64 GNU/Linux

```
[root@host1 ~]# ping -n 172.16.0.130
```

PING 172.16.0.130 (172.16.0.130) 56(84) bytes of data.

64 bytes from 172.16.0.130: icmp seq=1 ttl=64 time=522 ms

[illegible]

64 bytes from 172.16.0.130: icmp seq=2 ttl=64 time=0.149 ms

```
64 bytes from 172.16.0.130: icmp_seq=3 ttl=64 time=0.149 ms
```

```
***** host2:
```

```
[root@host2 ~]# iptables -t nat -L && iptables -t filter -L && iptables -t mangle -L
```

Chain PREROUTING (policy ACCEPT)

target	prot	opt	source	destination
1	1	1	1	1
2	2	2	2	2
3	3	3	3	3
4	4	4	4	4
5	5	5	5	5
6	6	6	6	6
7	7	7	7	7
8	8	8	8	8
9	9	9	9	9
10	10	10	10	10
11	11	11	11	11
12	12	12	12	12
13	13	13	13	13
14	14	14	14	14
15	15	15	15	15
16	16	16	16	16
17	17	17	17	17
18	18	18	18	18
19	19	19	19	19
20	20	20	20	20
21	21	21	21	21
22	22	22	22	22
23	23	23	23	23
24	24	24	24	24
25	25	25	25	25
26	26	26	26	26
27	27	27	27	27
28	28	28	28	28
29	29	29	29	29
30	30	30	30	30
31	31	31	31	31
32	32	32	32	32
33	33	33	33	33
34	34	34	34	34
35	35	35	35	35
36	36	36	36	36
37	37	37	37	37
38	38	38	38	38
39	39	39	39	39
40	40	40	40	40
41	41	41	41	41
42	42	42	42	42
43	43	43	43	43
44	44	44	44	44
45	45	45	45	45
46	46	46	46	46
47	47	47	47	47
48	48	48	48	48
49	49	49	49	49
50	50	50	50	50
51	51	51	51	51
52	52	52	52	52
53	53	53	53	53
54	54	54	54	54
55	55	55	55	55
56	56	56	56	56
57	57	57	57	57
58	58	58	58	58
59	59	59	59	59
60	60	60	60	60
61	61	61	61	61
62	62	62	62	62
63	63	63	63	63
64	64	64	64	64
65	65	65	65	65
66	66	66	66	66
67	67	67	67	67
68	68	68	68	68
69	69	69	69	69
70	70	70	70	70
71	71	71	71	71
72	72	72	72	72
73	73	73	73	73
74	74	74	74	74
75	75	75	75	75
76	76	76	76	76
77	77	77	77	77
78	78	78	78	78
79	79	79	79	79
80	80	80	80	80
81	81	81	81	81
82	82	82	82	82
83	83	83	83	83
84	84	84	84	84
85	85	85	85	85
86	86	86	86	86
87	87	87	87	87

Chain POSTROUTING (policy ACCEPT)

target	prot opt source		destination
--------	-----------------	--	-------------

Chain OUTPUT (policy ACCEPT)

target	prot opt	source	destination
--------	----------	--------	-------------

Chain INPUT (policy ACCEPT)

```
target    prot opt source      destination
```

Chain FORWARD (policy ACCEPT)

target	prot	opt	source	destination
--------	------	-----	--------	-------------

Chain OUTPUT (policy ACCEPT)
target prot opt source destination
Chain PREROUTING (policy ACCEPT)
target prot opt source destination

Chain INPUT (policy ACCEPT)
target prot opt source destination

Chain FORWARD (policy ACCEPT)
target prot opt source destination

Chain OUTPUT (policy ACCEPT)
target prot opt source destination

Chain POSTROUTING (policy ACCEPT)
target prot opt source destination

***** host1:

[root@host1 ~]# iptables -t nat -L && iptables -t filter -L && iptables -t mangle -L

Chain PREROUTING (policy ACCEPT)
target prot opt source destination

Chain POSTROUTING (policy ACCEPT)
target prot opt source destination

Chain OUTPUT (policy ACCEPT)
target prot opt source destination

Chain INPUT (policy ACCEPT)
target prot opt source destination

Chain FORWARD (policy ACCEPT)
target prot opt source destination

Chain OUTPUT (policy ACCEPT)
target prot opt source destination

Chain PREROUTING (policy ACCEPT)
target prot opt source destination

Chain INPUT (policy ACCEPT)
target prot opt source destination

Chain FORWARD (policy ACCEPT)
target prot opt source destination

Chain OUTPUT (policy ACCEPT)
target prot opt source destination

Chain POSTROUTING (policy ACCEPT)
target prot opt source destination

***** VE 1001:

```
[root@1001 /]# ip route list table all
192.0.2.0/24 dev venet0 scope host
169.254.0.0/16 dev venet0 scope link
default via 192.0.2.1 dev venet0
broadcast 127.255.255.255 dev lo table 255 proto kernel scope link src 127.0.0.1
local 172.16.0.130 dev venet0 table 255 proto kernel scope host src 172.16.0.130
broadcast 172.16.0.130 dev venet0 table 255 proto kernel scope link src 172.16.0.130
broadcast 127.0.0.0 dev lo table 255 proto kernel scope link src 127.0.0.1
local 127.0.0.1 dev lo table 255 proto kernel scope host src 127.0.0.1
local 127.0.0.1 dev venet0 table 255 proto kernel scope host src 127.0.0.1
local 127.0.0.0/8 dev lo table 255 proto kernel scope host src 127.0.0.1
unreachable default dev lo table unspec proto none metric -1 error -101 hoplimit 255
local ::1 via :: dev lo table 255 proto none metric 0 mtu 16436 advmss 16376 hoplimit
4294967295
unreachable default dev lo table unspec proto none metric -1 error -101 hoplimit 255
```

***** [root@host2 /]# ip route list table all

```
172.16.0.130 dev venet0 scope link
65.19.XXX.XX6/XX dev eth1 proto kernel scope link src 65.19.XXX.XX1
172.16.0.0/16 dev eth0 proto kernel scope link src 172.16.0.8
169.254.0.0/16 dev eth1 scope link
default via 65.19.XXX.XX7 dev eth1 src 65.19.XXX.XX1
broadcast 65.19.XXX.XX6 dev eth1 table 255 proto kernel scope link src 65.19.XXX.XX1
broadcast 127.255.255.255 dev lo table 255 proto kernel scope link src 127.0.0.1
local 65.19.XXX.XX1 dev eth1 table 255 proto kernel scope host src 65.19.XXX.XX1
broadcast 172.16.0.0 dev eth0 table 255 proto kernel scope link src 172.16.0.8
broadcast 172.16.255.255 dev eth0 table 255 proto kernel scope link src 172.16.0.8
local 172.16.0.8 dev eth0 table 255 proto kernel scope host src 172.16.0.8
broadcast 127.0.0.0 dev lo table 255 proto kernel scope link src 127.0.0.1
broadcast 65.19.XXX.X91 dev eth1 table 255 proto kernel scope link src 65.19.XXX.XX1
local 127.0.0.1 dev lo table 255 proto kernel scope host src 127.0.0.1
local 127.0.0.0/8 dev lo table 255 proto kernel scope host src 127.0.0.1
fe80::/64 dev eth0 metric 256 expires 16737768sec mtu 1500 advmss 1440 hoplimit
4294967295
fe80::/64 dev eth1 metric 256 expires 16737768sec mtu 1500 advmss 1440 hoplimit
4294967295
unreachable default dev lo table unspec proto none metric -1 error -101 hoplimit 255
local ::1 via :: dev lo table 255 proto none metric 0 mtu 16436 advmss 16376 hoplimit
4294967295
local fe80::230:48ff:fe8d:238a via :: dev lo table 255 proto none metric 0 mtu 16436 advmss
16376 hoplimit 4294967295
```

```
local fe80::230:48ff:fe8d:238b via :: dev lo table 255 proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
ff00::/8 dev eth0 table 255 metric 256 expires 16737768sec mtu 1500 advmss 1440 hoplimit 4294967295
ff00::/8 dev eth1 table 255 metric 256 expires 16737768sec mtu 1500 advmss 1440 hoplimit 4294967295
unreachable default dev lo table unspec proto none metric -1 error -101 hoplimit 255
```

```
[root@host2 conf]# tcpdump -i eth0 -n -e host 172.16.0.130
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 96 bytes
18:53:35.640769 00:25:90:34:25:28 > 00:30:48:8d:23:8a, ethertype IPv4 (0x0800), length 98: 172.16.0.20 > 172.16.0.130: ICMP echo request, id 15442, seq 1, length 64
18:53:35.640848 00:30:48:8d:23:8a > 00:25:90:34:25:28, ethertype IPv4 (0x0800), length 98: 172.16.0.130 > 172.16.0.20: ICMP echo reply, id 15442, seq 1, length 64
18:53:36.640122 00:25:90:34:25:28 > 00:30:48:8d:23:8a, ethertype IPv4 (0x0800), length 98: 172.16.0.20 > 172.16.0.130: ICMP echo request, id 15442, seq 2, length 64
18:53:36.640214 00:30:48:8d:23:8a > 00:25:90:34:25:28, ethertype IPv4 (0x0800), length 98: 172.16.0.130 > 172.16.0.20: ICMP echo reply, id 15442, seq 2, length 64
18:53:37.639852 00:25:90:34:25:28 > 00:30:48:8d:23:8a, ethertype IPv4 (0x0800), length 98: 172.16.0.20 > 172.16.0.130: ICMP echo request, id 15442, seq 3, length 64
18:53:37.639912 00:30:48:8d:23:8a > 00:25:90:34:25:28, ethertype IPv4 (0x0800), length 98: 172.16.0.130 > 172.16.0.20: ICMP echo reply, id 15442, seq 3, length 64
```

6 packets captured
6 packets received by filter

```
[root@host2 conf]# tcpdump -i venet0 -n -e host 172.16.0.130
tcpdump: WARNING: arptype 65535 not supported by libpcap - falling back to cooked socket
tcpdump: WARNING: venet0: no IPv4 address assigned
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on venet0, link-type LINUX_SLL (Linux cooked), capture size 96 bytes
18:53:10.768100 Out ethertype IPv4 (0x0800), length 100: 172.16.0.20 > 172.16.0.130: ICMP echo request, id 47441, seq 1, length 64
18:53:10.768146 In ethertype IPv4 (0x0800), length 100: 172.16.0.130 > 172.16.0.20: ICMP echo reply, id 47441, seq 1, length 64
18:53:11.767602 Out ethertype IPv4 (0x0800), length 100: 172.16.0.20 > 172.16.0.130: ICMP echo request, id 47441, seq 2, length 64
18:53:11.767650 In ethertype IPv4 (0x0800), length 100: 172.16.0.130 > 172.16.0.20: ICMP echo reply, id 47441, seq 2, length 64
18:53:12.767384 Out ethertype IPv4 (0x0800), length 100: 172.16.0.20 > 172.16.0.130: ICMP echo request, id 47441, seq 3, length 64
18:53:12.767433 In ethertype IPv4 (0x0800), length 100: 172.16.0.130 > 172.16.0.20: ICMP echo reply, id 47441, seq 3, length 64
```

6 packets captured

6 packets received by filter
0 packets dropped by kernel
