
Subject: Re: Source based routing

Posted by [sspt](#) on Sat, 29 Mar 2008 03:31:48 GMT

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

It happens from any external box which tries to reach the VE

External BOX

ping to VE

```
21:50:09.342794 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 0
21:50:10.351723 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 1
21:50:11.360495 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 2
21:50:12.360282 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 3
21:50:13.360058 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 4
21:50:14.359834 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 5
21:50:15.359627 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 6
21:50:16.368357 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 7
21:50:17.368154 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 8
21:50:18.367927 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 9
21:50:19.367700 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 10
21:50:20.367476 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 11
21:50:21.376229 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 12
21:50:22.376023 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 13
21:50:23.379326 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.140: icmp 64: echo request seq 14
```

ping to node

```
21:59:08.040177 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.130: icmp 64: echo request seq 0
21:59:08.113576 IP (tos 0x0, ttl 40, id 56698, offset 0, flags [none], proto 1, length: 84)
67.xxx.xxx.130 > xxx.110.83.42: icmp 64: echo reply seq 0
21:59:09.040162 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.130: icmp 64: echo request seq 1
21:59:09.113478 IP (tos 0x0, ttl 40, id 56699, offset 0, flags [none], proto 1, length: 84)
67.xxx.xxx.130 > xxx.110.83.42: icmp 64: echo reply seq 1
21:59:10.039948 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
```

```

67.xxx.xxx.130: icmp 64: echo request seq 2
21:59:10.113954 IP (tos 0x0, ttl 40, id 56700, offset 0, flags [none], proto 1, length: 84)
67.xxx.xxx.130 > xxx.110.83.42: icmp 64: echo reply seq 2
21:59:11.039702 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.130: icmp 64: echo request seq 3
21:59:11.113510 IP (tos 0x0, ttl 40, id 56701, offset 0, flags [none], proto 1, length: 84)
67.xxx.xxx.130 > xxx.110.83.42: icmp 64: echo reply seq 3
21:59:12.048458 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.130: icmp 64: echo request seq 4
21:59:12.121922 IP (tos 0x0, ttl 40, id 56702, offset 0, flags [none], proto 1, length: 84)
67.xxx.xxx.130 > xxx.110.83.42: icmp 64: echo reply seq 4
21:59:13.057232 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.130: icmp 64: echo request seq 5
21:59:13.130931 IP (tos 0x0, ttl 40, id 56703, offset 0, flags [none], proto 1, length: 84)
67.xxx.xxx.130 > xxx.110.83.42: icmp 64: echo reply seq 5
21:59:14.057046 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.130: icmp 64: echo request seq 6
21:59:14.129665 IP (tos 0x0, ttl 40, id 56704, offset 0, flags [none], proto 1, length: 84)
67.xxx.xxx.130 > xxx.110.83.42: icmp 64: echo reply seq 6
21:59:15.056793 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.130: icmp 64: echo request seq 7
21:59:15.129523 IP (tos 0x0, ttl 40, id 56705, offset 0, flags [none], proto 1, length: 84)
67.xxx.xxx.130 > xxx.110.83.42: icmp 64: echo reply seq 7
21:59:16.056568 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.130: icmp 64: echo request seq 8
21:59:16.141422 IP (tos 0x0, ttl 40, id 56706, offset 0, flags [none], proto 1, length: 84)
67.xxx.xxx.130 > xxx.110.83.42: icmp 64: echo reply seq 8
21:59:17.056378 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.130: icmp 64: echo request seq 9
21:59:17.129063 IP (tos 0x0, ttl 40, id 56707, offset 0, flags [none], proto 1, length: 84)
67.xxx.xxx.130 > xxx.110.83.42: icmp 64: echo reply seq 9
21:59:18.056139 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) xxx.110.83.42 >
67.xxx.xxx.130: icmp 64: echo request seq 10
21:59:18.128998 IP (tos 0x0, ttl 40, id 56708, offset 0, flags [none], proto 1, length: 84)
67.xxx.xxx.130 > xxx.110.83.42: icmp 64: echo reply seq 10

```

Node

While we ping VE IP

Nothing, a traceroute won't reach the VE IP, it dies at the switch

So i've ran '/usr/sbin/arpsh -U -i 67.xxx.xxx.140 -c 1 eth1' and then:

```

20:20:07.121016 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.140: ICMP echo request, id 60987, seq 81, length 64
20:20:07.121108 IP (tos 0x0, ttl 64, id 56946, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.140 > xxx.110.83.42: ICMP echo reply, id 60987, seq 81, length 64
20:20:07.121116 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.140: ICMP echo request, id 60987, seq 82, length 64
20:20:07.121142 IP (tos 0x0, ttl 64, id 56947, offset 0, flags [none], proto: ICMP (1), length: 84)

```

67.xxx.xxx.140 > xxx.110.83.42: ICMP echo reply, id 60987, seq 82, length 64
20:20:07.121643 IP (tos 0xc0, ttl 64, id 6287, offset 0, flags [none], proto: ICMP (1), length: 576)
67.xxx.xxx.140 > 198.32.64.12: ICMP 67.xxx.xxx.140 udp port 32915 unreachable, length 556
20:20:07.121763 IP (tos 0xc0, ttl 64, id 2440, offset 0, flags [none], proto: ICMP (1), length: 163)
67.xxx.xxx.140 > 192.228.79.201: ICMP 67.xxx.xxx.140 udp port 32915 unreachable, length 143
20:20:07.122048 IP (tos 0xc0, ttl 64, id 641, offset 0, flags [none], proto: ICMP (1), length: 341)
67.xxx.xxx.140 > 192.203.230.10: ICMP 67.xxx.xxx.140 udp port 32915 unreachable, length 321
20:20:08.052997 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.140: ICMP echo request, id 63035, seq 8, length 64
20:20:08.053030 IP (tos 0x0, ttl 64, id 56948, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.140 > xxx.110.83.42: ICMP echo reply, id 63035, seq 8, length 64
20:20:09.048410 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.140: ICMP echo request, id 63035, seq 9, length 64
20:20:09.048439 IP (tos 0x0, ttl 64, id 56949, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.140 > xxx.110.83.42: ICMP echo reply, id 63035, seq 9, length 64
20:20:10.047784 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.140: ICMP echo request, id 63035, seq 10, length 64
20:20:10.047815 IP (tos 0x0, ttl 64, id 56950, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.140 > xxx.110.83.42: ICMP echo reply, id 63035, seq 10, length 64
20:20:11.047685 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.140: ICMP echo request, id 63035, seq 11, length 64
20:20:11.047713 IP (tos 0x0, ttl 64, id 56951, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.140 > xxx.110.83.42: ICMP echo reply, id 63035, seq 11, length 64
20:20:12.047350 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.140: ICMP echo request, id 63035, seq 12, length 64
20:20:12.047379 IP (tos 0x0, ttl 64, id 56952, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.140 > xxx.110.83.42: ICMP echo reply, id 63035, seq 12, length 64
20:20:13.047414 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.140: ICMP echo request, id 63035, seq 13, length 64
20:20:13.047442 IP (tos 0x0, ttl 64, id 56953, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.140 > xxx.110.83.42: ICMP echo reply, id 63035, seq 13, length 64
20:20:14.047377 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.140: ICMP echo request, id 63035, seq 14, length 64
20:20:14.047406 IP (tos 0x0, ttl 64, id 56954, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.140 > xxx.110.83.42: ICMP echo reply, id 63035, seq 14, length 64
20:20:15.048324 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.140: ICMP echo request, id 63035, seq 15, length 64
20:20:15.048352 IP (tos 0x0, ttl 64, id 56955, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.140 > xxx.110.83.42: ICMP echo reply, id 63035, seq 15, length 64
20:20:16.046954 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.140: ICMP echo request, id 63035, seq 16, length 64
A few minutes later, nothing again

While we ping node IP

20:17:43.646730 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.130: ICMP echo request, id 20283, seq 0, length 64
20:17:43.646757 IP (tos 0x0, ttl 64, id 56897, offset 0, flags [none], proto: ICMP (1), length: 84)

```

67.xxx.xxx.130 > xxx.110.83.42: ICMP echo reply, id 20283, seq 0, length 64
20:17:44.647210 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.130: ICMP echo request, id 20283, seq 1, length 64
20:17:44.647225 IP (tos 0x0, ttl 64, id 56898, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.130 > xxx.110.83.42: ICMP echo reply, id 20283, seq 1, length 64
20:17:45.647314 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.130: ICMP echo request, id 20283, seq 2, length 64
20:17:45.647327 IP (tos 0x0, ttl 64, id 56899, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.130 > xxx.110.83.42: ICMP echo reply, id 20283, seq 2, length 64
20:17:46.773103 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.130: ICMP echo request, id 20283, seq 3, length 64
20:17:46.773115 IP (tos 0x0, ttl 64, id 56900, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.130 > xxx.110.83.42: ICMP echo reply, id 20283, seq 3, length 64
20:17:47.647223 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.130: ICMP echo request, id 20283, seq 4, length 64
20:17:47.647237 IP (tos 0x0, ttl 64, id 56901, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.130 > xxx.110.83.42: ICMP echo reply, id 20283, seq 4, length 64
20:17:48.647086 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.130: ICMP echo request, id 20283, seq 5, length 64
20:17:48.647099 IP (tos 0x0, ttl 64, id 56902, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.130 > xxx.110.83.42: ICMP echo reply, id 20283, seq 5, length 64
20:17:49.647207 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.130: ICMP echo request, id 20283, seq 6, length 64
20:17:49.647221 IP (tos 0x0, ttl 64, id 56903, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.130 > xxx.110.83.42: ICMP echo reply, id 20283, seq 6, length 64
20:17:50.647366 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.130: ICMP echo request, id 20283, seq 7, length 64
20:17:50.647380 IP (tos 0x0, ttl 64, id 56904, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.130 > xxx.110.83.42: ICMP echo reply, id 20283, seq 7, length 64
20:17:51.647339 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.130: ICMP echo request, id 20283, seq 8, length 64
20:17:51.647355 IP (tos 0x0, ttl 64, id 56905, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.130 > xxx.110.83.42: ICMP echo reply, id 20283, seq 8, length 64
20:17:52.648187 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.130: ICMP echo request, id 20283, seq 9, length 64
20:17:52.648199 IP (tos 0x0, ttl 64, id 56906, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.130 > xxx.110.83.42: ICMP echo reply, id 20283, seq 9, length 64
20:17:53.647165 IP (tos 0x0, ttl 40, id 0, offset 0, flags [DF], proto: ICMP (1), length: 84)
xxx.110.83.42 > 67.xxx.xxx.130: ICMP echo request, id 20283, seq 10, length 64
20:17:53.647177 IP (tos 0x0, ttl 64, id 56907, offset 0, flags [none], proto: ICMP (1), length: 84)
67.xxx.xxx.130 > xxx.110.83.42: ICMP echo reply, id 20283, seq 10, length 64

```

```

[root@node ~]# ip route | table 7
default via 67.xxx.xxx.129 dev eth1

```

```

[root@node ~]# ip route | table 8

```

default via 69.xxx.xx0.129 dev eth0

```
[root@node ~]# ip route | table main
67.xxx.xxx.140 dev venet0 scope link
69.xxx.xx0.179 dev venet0 scope link
69.xxx.xx0.187 dev venet0 scope link
67.xxx.xxx.128/28 dev eth1 proto kernel scope link src 67.xxx.xxx.130
69.xxx.xx0.128/26 dev eth0 proto kernel scope link src 69.xxx.xx0.181
169.254.0.0/16 dev eth1 scope link
default via 69.xxx.xx0.129 dev eth0
```

```
[root@node ~]# ip route | table 255
broadcast 127.255.255.255 dev lo proto kernel scope link src 127.0.0.1
local 69.xxx.xx0.181 dev eth0 proto kernel scope host src 69.xxx.xx0.181
broadcast 67.xxx.xxx.142 dev eth1 proto kernel scope link src 67.xxx.xxx.130
broadcast 67.xxx.xxx.143 dev eth1 proto kernel scope link src 67.xxx.xxx.130
broadcast 69.xxx.xx0.128 dev eth0 proto kernel scope link src 69.xxx.xx0.181
broadcast 67.xxx.xxx.128 dev eth1 proto kernel scope link src 67.xxx.xxx.130
broadcast 69.xxx.xx0.191 dev eth0 proto kernel scope link src 69.xxx.xx0.181
local 67.xxx.xxx.130 dev eth1 proto kernel scope host src 67.xxx.xxx.130
broadcast 127.0.0.0 dev lo proto kernel scope link src 127.0.0.1
local 127.0.0.1 dev lo proto kernel scope host src 127.0.0.1
local 127.0.0.0/8 dev lo proto kernel scope host src 127.0.0.1
```

```
[root@node ~]# sysctl -a | grep rp_filter
net.ipv4.conf.venet0.arp_filter = 0
net.ipv4.conf.venet0.rp_filter = 1
net.ipv4.conf.eth1.arp_filter = 0
net.ipv4.conf.eth1.rp_filter = 1
net.ipv4.conf.eth0.arp_filter = 0
net.ipv4.conf.eth0.rp_filter = 1
net.ipv4.conf.lo.arp_filter = 0
net.ipv4.conf.lo.rp_filter = 0
net.ipv4.conf.default.arp_filter = 0
net.ipv4.conf.default.rp_filter = 1
net.ipv4.conf.all.arp_filter = 0
net.ipv4.conf.all.rp_filter = 0
```

Thank you for all the support
