Subject: Networking problem (internet is unreachable)
Posted by RoverCZ on Sat, 06 Jun 2009 20:42:30 GMT
View Forum Message <> Reply to Message

Hello,

I have a problem with network traffic routing probably. I've set up two HW nodes, both have public IPv4 address. I'm running about 30 openvz containers on each of them. All the containers were working fine until yesterday. I've noticed that any another added VE on HW1 cannot communicate outside the HW node. All the existing VE are working fine. If I try to add the VE on HW2 exactly same way, it works ok.

I've tried to capture packets with tcpdump pinging from HW2 to "not working" VE. Ping (nor TCP/UDP communication) does not work from any other computer either (but it was easier to monitor it on linux).

HW2 (sender of ping command):

hw02:~# tcpdump -n -i eth0 -e host 88.86.119.76
tcpdump: WARNING: eth0: no IPv4 address assigned
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 96 bytes
22:23:01.590038 00:30:48:c5:39:06 > ff:ff:ff:ff:ff:ff, ethertype ARP (0x0806), length 42: arp who-has 88.86.119.76 tell 88.86.119.2
22:23:02.253020 00:15:17:77:80:e4 > 00:30:48:c5:39:06, ethertype ARP (0x0806), length 60: arp reply 88.86.119.76 is-at 00:15:17:77:80:e4
22:23:02.253044 00:30:48:c5:39:06 > 00:15:17:77:80:e4, ethertype IPv4 (0x0800), length 98: 88.86.119.2 > 88.86.119.76: ICMP echo request, id 29030, seq 1, length 64

HW1 (where the VE is on), eth0:

hw01:~# tcpdump -n -i eth0 -e host 88.86.119.76
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 96 bytes
22:23:01.590028 00:30:48:c5:39:06 > ff:ff:ff:ff:ff:ff, ethertype ARP (0x0806), length 60: arp who-has 88.86.119.76 tell 88.86.119.2
22:23:02.252770 00:15:17:77:80:e4 > 00:30:48:c5:39:06, ethertype ARP (0x0806), length 42: arp reply 88.86.119.76 is-at 00:15:17:77:80:e4
22:23:02.252915 00:30:48:c5:39:06 > 00:15:17:77:80:e4, ethertype IPv4 (0x0800), length 98: 88.86.119.2 > 88.86.119.76: ICMP echo request, id 29030, seq 1, length 64
22:23:02.253027 00:15:17:77:80:e4 > 00:30:48:c5:39:06, ethertype IPv4 (0x0800), length 98: 88.86.119.76 > 88.86.119.2: ICMP echo reply, id 29030, seq 1, length 64

HW1 venet0:

```
hw01:~# tcpdump -n -i venet0 -e host 88.86.119.76
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
22:23:02.252942 Out ethertype IPv4 (0x0800), length 100: 88.86.119.2 > 88.86.119.76: ICMP
echo request, id 29030, seq 1, length 64
22:23:02.252994  In ethertype IPv4 (0x0800), length 100: 88.86.119.76 > 88.86.119.2: ICMP echo
reply, id 29030, seq 1, length 64
```

VE:

```
vs75:/# tcpdump -n
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
22:23:02.252942 IP 88.86.119.2 > 88.86.119.76: ICMP echo request, id 29030, seq 1, length 64
22:23:02.252994 IP 88.86.119.76 > 88.86.119.2: ICMP echo reply, id 29030, seq 1, length 64
```

HW1:

```
Linux hw01 2.6.18-14-fza-amd64 #1 SMP Mon Jan 5 17:36:46 UTC 2009 x86_64 GNU/Linux
```

IPTables are empty

```
hw01:~# ip rule list
0:      from all lookup 255
32766:  from all lookup main
32767:  from all lookup default
```

```
hw01:~# cat /etc/sysctl.conf
#
# /etc/sysctl.conf - Configuration file for setting system variables
# See sysctl.conf (5) for information.
#

#kernel.domainname = example.com
#net/ipv4/icmp_echo_ignore_broadcasts=1

# Uncomment the following to stop low-level messages on console
#kernel.printk = 4 4 1 7

##############################################################3
# Functions previously found in netbase
#
```

```
# Uncomment the next line to enable Spoof protection (reverse-path filter)
#net.ipv4.conf.default.rp_filter=1

# Uncomment the next line to enable TCP/IP SYN cookies
#net.ipv4.tcp_syncookies=1

# Uncomment the next line to enable packet forwarding for IPv4
#net.ipv4.conf.default.forwarding=1

# Uncomment the next line to enable packet forwarding for IPv6
#net.ipv6.conf.default.forwarding=1

#-- OpenVZ begin --#

# On Hardware Node we generally need
# packet forwarding enabled and proxy arp disabled
net.ipv4.conf.default.forwarding=1
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.ip_forward=1

# Enables source route verification
net.ipv4.conf.all.rp_filter = 1

# Enables the magic-sysrq key
kernel.sysrq = 1

# TCP Explict Congestion Notification
#net.ipv4.tcp_ecn = 0

# we do not want all our interfaces to send redirects
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0

#-- OpenVZ end --#
```

HW2:
Linux hw02 2.6.24-2-pve #1 SMP PREEMPT Wed Jan 14 11:32:49 CET 2009 x86_64 GNU/Linux
IPTables are empty

```
hw02:~# ip rule list
0:      from all lookup 255
32766:  from all lookup main
32767:  from all lookup default


hw02:~# cat /etc/sysctl.conf
#
```

```
# /etc/sysctl.conf - Configuration file for setting system variables
# See sysctl.conf (5) for information.
#

#kernel.domainname = example.com
#net/ipv4/icmp_echo_ignore_broadcasts=1

# Uncomment the following to stop low-level messages on console
#kernel.printk = 4 4 1 7

######################################################################3
# Functions previously found in netbase
#

# Uncomment the next line to enable Spoof protection (reverse-path filter)
#net.ipv4.conf.default.rp_filter=1

# Uncomment the next line to enable TCP/IP SYN cookies
#net.ipv4.tcp_syncookies=1

# Uncomment the next line to enable packet forwarding for IPv4
#net.ipv4.conf.default.forwarding=1

# Uncomment the next line to enable packet forwarding for IPv6
#net.ipv6.conf.default.forwarding=1

#-- OpenVZ begin --#

# On Hardware Node we generally need
# packet forwarding enabled and proxy arp disabled
net.ipv4.conf.default.forwarding=1
net.ipv4.conf.default.proxy_arp = 0

# Enables source route verification
net.ipv4.conf.all.rp_filter = 1

# Enables the magic-sysrq key
kernel.sysrq = 1

# TCP Explict Congestion Notification
#net.ipv4.tcp_ecn = 0

# we do not want all our interfaces to send redirects
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0

#-- OpenVZ end --#
```

I just can't figure out where the problem is, can you help me please?