Subject: IPv6 again but this time with veth
Posted by Jan Tomasek on Wed, 27 Jun 2007 15:42:06 GMT
View Forum Message <> Reply to Message

Hello,

few weeks ago I was trying to setup IPv6, but that doesn't work.
http://forum.openvz.org/index.php?t=msg&goto=12809&& amp;srch=IPv6#msg_12809

I would really like to have this feature. So I'm trying to experiment
with veth devices, but not much success again :(

My host node has IPv6 support:

```
> staj-dev:~# ifconfig
> eth0      Link encap:Ethernet  HWaddr 00:13:72:1B:0B:97
>           inet addr:195.113.233.4  Bcast:195.113.233.255  Mask:255.255.255.0
>           inet6 addr: 2001:718:1:e::23:3004/64 Scope:Global
>           inet6 addr: fe80::213:72ff:fe1b:b97/64 Scope:Link
>           UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
>           RX packets:2136 errors:0 dropped:0 overruns:0 frame:0
>           TX packets:1304 errors:0 dropped:0 overruns:0 carrier:0
>           collisions:0 txqueuelen:1000
>           RX bytes:211656 (206.6 KiB)  TX bytes:138480 (135.2 KiB)
>           Interrupt:17
```

It is accessible on IPv6:

```
 > semik@staff:~$ ping6 -c1 2001:718:1:e::23:3004
 > PING 2001:718:1:e::23:3004(2001:718:1:e::23:3004) 56 data bytes
 > 64 bytes from 2001:718:1:e::23:3004: icmp_seq=0 ttl=63 time=0.129 ms
```

I setup IPv4 part this way:

```
 > staj-dev:~# ifconfig veth012.0 0
 > staj-dev:~# echo 1 > /proc/sys/net/ipv4/conf/veth012.0/forwarding
 > staj-dev:~# echo 1 > /proc/sys/net/ipv4/conf/veth012.0/proxy_arp
 > staj-dev:~# echo 1 > /proc/sys/net/ipv4/conf/eth0/forwarding
 > staj-dev:~# echo 1 > /proc/sys/net/ipv4/conf/eth0/proxy_arp
 > staj-dev:~# ip route add 195.113.233.012 dev veth012.0
```

and it works:

```
 > semik:~$ping -c1 195.113.233.12
 > PING 195.113.233.12 (195.113.233.12): 56 data bytes
 > 64 bytes from 195.113.233.12: icmp_seq=0 ttl=61 time=0.252 ms
```

I'm trying to follow

http://wiki.openvz.org/Virtual_Ethernet_device#Virtual_ether net_device_with_IPv6

but after executing (2nd, 3rd) lines

> staj-dev:~# echo 1 > /proc/sys/net/ipv6/conf/veth012.0/forwarding
> staj-dev:~# echo 1 > /proc/sys/net/ipv6/conf/eth0/forwarding
> staj-dev:~# echo 1 > /proc/sys/net/ipv6/conf/all/forwarding

host node stop responding to ping6. I'm stuck here... I also do not want
run router advertisement daemon (if that is possible). My VE should have
address 2001:718:1:e::23:3012.

Inside VE I setup:

> root@ip6v2-test:~# ifconfig
> eth0 Link encap:Ethernet HWaddr 00:0C:29:63:A6:73
> inet addr:195.113.233.12 Bcast:195.113.233.255 Mask:255.255.255.0
> inet6 addr: 2001:718:1:e::23:3012/64 Scope:Global
> inet6 addr: fe80::20c:29ff:fe63:a673/64 Scope:Link
> UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
> RX packets:88 errors:0 dropped:0 overruns:0 frame:0
> TX packets:88 errors:0 dropped:0 overruns:0 carrier:0
> collisions:0 txqueuelen:0
> RX bytes:7012 (6.8 KiB) TX bytes:6948 (6.7 KiB)

>
> root@ip6v2-test:~# cat /etc/network/interfaces
> auto lo
> iface lo inet loopback
>
> auto eth0
> iface eth0 inet static
> address 195.113.233.12
> netmask 255.255.255.0
> network 195.113.233.0
> broadcast 195.113.233.255
> gateway 195.113.233.4
> dns-nameservers 195.113.144.233 195.113.144.194
> dns-search cesnet.cz
>
>
> eth0 inet6 static
> iface eth0 inet6 static
> address 2001:718:1:e::23:3012
> netmask 64

veth interface is configured this way:

> NETIF=" ifname=eth0,mac=00:0C:29:63:A6:73,host_ifname=veth012.0,host
_mac=00:0C:29:63:A6:74 "

inside VE I see reasonable routes:

> root@ip6v2-test:~# route -6
> Kernel IPv6 routing table
> Destination Next Hop Flags Metric Ref Use Iface
> ::1/128 :: U 0 1 1 lo
> 2001:718:1:e::23:3004/128 2001:718:1:e::23:3004 UC 0 49 0
eth0
> 2001:718:1:e::23:3012/128 :: U 0 142 1 lo
> 2001:718:1:e::/64 :: U 256 0 0 eth0
> fe80::20c:29ff:fe63:a673/128 :: U 0 0 1 lo
> fe80::/64 :: U 256 0 0 eth0
> ff02::16/128 ff02::16 UC 0 5 0 eth0
> ff00::/8 :: U 256 0 0 eth0

but I'm not able even to ping to my host:

> root@ip6v2-test:~# ping6 2001:718:1:e::23:3004
> PING 2001:718:1:e::23:3004(2001:718:1:e::23:3004) 56 data bytes
> From 2001:718:1:e::23:3012 icmp_seq=1 Destination unreachable: Address unreachable

it is very interesting, that host node see my ping requests:

> staj-dev:~# tcpdump -i veth012.0
> tcpdump: WARNING: veth012.0: no IPv4 address assigned
> tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
> listening on veth012.0, link-type EN10MB (Ethernet), capture size 96 bytes
> 18:48:53.301938 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation,
who has staj-dev6.cesnet.cz, length 32
> 18:48:54.301800 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation,
who has staj-dev6.cesnet.cz, length 32
> 18:48:55.302659 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation,
who has staj-dev6.cesnet.cz, length 32

but do not respond to ping :(

I'm trying this on 2.6.18-028stab035 system with vzctl version 3.0.16-5dso1.

Could anybody help? Thanks.

--

Jan Tomasek aka Semik
http://www.tomasek.cz/

Subject: Re: IPv6 again but this time with veth
Posted by Jan Tomasek on Thu, 28 Jun 2007 14:50:27 GMT
View Forum Message <> Reply to Message

Vitaliy Gusev wrote:
> Hello, do you use a bridge?

Probably not... if there is some place where I should RTFM, than just provide me URL please ;) I was hoping setup will be similar as is for IPv4. IPv6 itself is new for me, and with combination with OpenVZ quite hard to understand

> How I see <staj-dev> is a host-node,

Yes.

> but what is two others hosts <semik>,

that is just IPv4 capable box, used for demonstration that my veth012.0 interface is at least able to communicate on IPv4.

 > <ip6v2-test>?

That is VE 233012 where I'm trying to get IPv6 working.

> What steps did you do for veth012.0 ?

staj-dev# vzctl set 233012 --netif_add eth0,00:0C:29:63:A6:73,veth012.0,00:0C:29:63:A6:74 --save

>> > staj-dev:~# ifconfig veth012.0 0
>> > staj-dev:~# echo 1 > /proc/sys/net/ipv4/conf/veth012.0/forwarding
>> > staj-dev:~# echo 1 > /proc/sys/net/ipv4/conf/veth012.0/proxy_arp
>> > staj-dev:~# echo 1 > /proc/sys/net/ipv4/conf/eth0/forwarding
>> > staj-dev:~# echo 1 > /proc/sys/net/ipv4/conf/eth0/proxy_arp
>> > staj-dev:~# ip route add 195.113.233.012 dev veth012.0
>>
>> > staj-dev:~# echo 1 > /proc/sys/net/ipv6/conf/veth012.0/forwarding
>> > staj-dev:~# echo 1 > /proc/sys/net/ipv6/conf/eth0/forwarding
>> > staj-dev:~# echo 1 > /proc/sys/net/ipv6/conf/all/forwarding

Am I now more clear? Thanks for help.

--

Jan Tomasek aka Semik
http://www.tomasek.cz/

Subject: Re: IPv6 again but this time with veth
Posted by Jan Tomasek on Thu, 28 Jun 2007 18:44:48 GMT
View Forum Message <> Reply to Message

Vitaliy Gusev wrote:
> Well, can you draw interaction scheme of the all hosts and describe final goal that you expect?

I need IPv4 and IPv6 support in VE.

Host node have IPv4=195.113.233.4, IPv6=2001:718:1:e::23:3004/64. That
is working ok.

VE 233012 is supposed to have IPv4=195.113.233.12,
IPv6=2001:718:1:e::23:3012/64.

Gateway for IPv4 is 195.113.233.1 and for IPv6 it somehow works without
GW. Config from host node:

> auto eth0
> iface eth0 inet static
> address 195.113.233.4
> netmask 255.255.255.0
> network 195.113.233.0
> broadcast 195.113.233.255
> gateway 195.113.233.1
> dns-nameservers 195.113.144.233 195.113.144.194
> dns-search cesnet.cz
>
> iface eth0 inet6 static
> post-up /sbin/sysctl -w "net.ipv6.conf.eth0.autoconf=0"
> address 2001:718:1:e::23:3004
> netmask 64

I've no idea how to draw that. Task is simple: Get IPv6 support inside
VE 233012.

I would prefer venet interface, but that doesn't work. See thread:
 http://forum.openvz.org/index.php?t=msg&goto=12809&& amp;srch=IPv6#msg_12809

So I'm trying experiment with veth.

> I want to make a note that veth device is usually used with a bridge.

Ok. I will use anything which will make it to work. I tried to restart
staj-dev to get clean environment and typed:

staj-dev:~# brctl addbr vzbr0
staj-dev:~# brctl addif vzbr0 veth012.0

staj-dev:~# echo 1 > /proc/sys/net/ipv4/conf/vzbr0/forwarding
staj-dev:~# echo 1 > /proc/sys/net/ipv4/conf/vzbr0/proxy_arp
staj-dev:~# ip route add 195.113.233.12 dev vzbr0
staj-dev:~# echo 1 > /proc/sys/net/ipv4/conf/eth0/forwarding
staj-dev:~# echo 1 > /proc/sys/net/ipv4/conf/eth0/proxy_arp
staj-dev:~# ip route add 2001:718:1:e::23:3012/64 dev vzbr0
staj-dev:~# ip route add 2001:718:1:e::23:3012 dev vzbr0
staj-dev:~# echo 1 > /proc/sys/net/ipv6/conf/eth0/forwarding
staj-dev:~# echo 1 > /proc/sys/net/ipv6/conf/vzbr0/forwarding

IPv4 is working fine. But IPv6 still not. :(

Thank you for your help so far.

Do you have any other suggestions?
--

Jan Tomasek aka Semik
http://www.tomasek.cz/

Subject: Re: IPv6 again but this time with veth
Posted by Jan Tomasek on Fri, 29 Jun 2007 07:35:47 GMT
View Forum Message <> Reply to Message

Vitaliy Gusev wrote:
> I still have some questions.
>
> You wrote that the host <semik> (or <staff>) is an IPv4 capable box,
>
>>> but what is two others hosts <semik>,
>> that is just IPv4 capable box, used for demonstration that my veth012.0
>> interface is at least able to communicate on IPv4.
>
> but I see you also use it as an IPv6 host:

<semik> is capable only of IPv4 (it is my workstation, that is reason
why mention it here, sorry for complicating things).

IPv6 is available on <staff> I used it for testing if I correctly setup
IPv6 on my host node.

>> semik@staff:~$ ping6 -c1 2001:718:1:e::23:3004
> > PING 2001:718:1:e::23:3004(2001:718:1:e::23:3004) 56 data bytes
> > 64 bytes from 2001:718:1:e::23:3004: icmp_seq=0 ttl=63 time=0.129 ms
>
> In other words IPV6 worked but you did something and those pings was stopped.

>
> Did I understand right?

No IPv6 worked and works for host node IPv6=2001:718:1:e::23:3004.

It never worked for VE 233012 with IPv6=2001:718:1:e::23:3012.

--

Jan Tomasek aka Semik
http://www.tomasek.cz/

Subject: Re: IPv6 again but this time with veth
Posted by Jan Tomasek on Fri, 29 Jun 2007 08:23:28 GMT
View Forum Message <> Reply to Message

Oh. I got lost! Could I try to describe it again?

My HW node is staj-dev:

```
auto eth0
iface eth0 inet static
        address 195.113.233.4
        netmask 255.255.255.0
        network 195.113.233.0
        broadcast 195.113.233.255
        gateway 195.113.233.1
        dns-nameservers 195.113.144.233 195.113.144.194
        dns-search cesnet.cz

iface eth0 inet6 static
        post-up /sbin/sysctl -w "net.ipv6.conf.eth0.autoconf=0"
        address 2001:718:1:e::23:3004
        netmask 64
```

My VE node, has this config (ip6v2-test):

```
auto eth0
iface eth0 inet static
        address 195.113.233.12
        netmask 255.255.255.0
        network 195.113.233.0
        broadcast 195.113.233.255
        gateway 195.113.233.4
        dns-nameservers 195.113.144.233 195.113.144.194
        dns-search cesnet.cz
```

eth0 inet6 static
        iface eth0 inet6 static
        address 2001:718:1:e::23:3012
        netmask 64

After executing:

brctl addbr vzbr0
brctl addif vzbr0 veth012.0
ifconfig vzbr0 0
echo 1 > /proc/sys/net/ipv4/conf/vzbr0/forwarding
echo 1 > /proc/sys/net/ipv4/conf/vzbr0/proxy_arp
echo 1 > /proc/sys/net/ipv4/conf/eth0/forwarding
echo 1 > /proc/sys/net/ipv4/conf/eth0/proxy_arp
ip route add 195.113.233.12 dev vzbr0
ip route add 2001:718:1:e::23:3012 dev vzbr0

MARK

VE starts to be reachable from Internet on IPv4. What should I do next?

echo 1 > /proc/sys/net/ipv6/conf/eth0/forwarding

somehow broke IPv6 for HW node.

echo 1 > /proc/sys/net/ipv6/conf/vzbr0/forwarding

cause no problem.

Well - returning to point marked as **MARK**.

Route table on HW node:

> staj-dev:~# route -6
> Kernel IPv6 routing table
> Destination Next Hop Flags Metric Ref Use Iface
> ::1/128 :: U 0 6 1 lo
> 2001:718:1:e::23:3004/128 :: U 0 788 1 lo
> 2001:718:1:e::23:3012/128 :: U 1024 0 0 vzbr0
> 2001:718:1:e::/64 :: U 256 169 0 eth0
> fe80::/128 :: U 0 0 2 lo
> fe80::1/128 :: U 0 0 1 lo
> fe80::20c:29ff:fe63:a674/128 :: U 0 0 1 lo
> fe80::20c:29ff:fe63:a674/128 :: U 0 0 1 lo
> fe80::213:72ff:fe1b:b97/128 :: U 0 32 1 lo
> fe80::/64 :: U 256 0 0 eth0
> fe80::/64 :: U 256 0 0 venet0

> fe80::/64 :: U 256 0 0 veth012.0
> fe80::/64 :: U 256 0 0 vzbr0
> f020::1/128 ff02::1 UC 0 165 0 eth0
> ff00::/8 :: U 256 0 0 eth0
> ff00::/8 :: U 256 0 0 venet0
> ff00::/8 :: U 256 0 0 veth012.0
> ff00::/8 :: U 256 0 0 vzbr0
> ::/0 fe80::215:faff:fe87:3100 UGDA 1024 167 0 eth0

Route table on VE node:

> root@ip6v2-test:/# route -6
> Kernel IPv6 routing table
> Destination Next Hop Flags Metric Ref Use Iface
> ::1/128 :: U 0 1 1 lo
> 2001:718:1:e::23:3012/128 :: U 0 53 1 lo
> 2001:718:1:e::/64 :: U 256 0 0 eth0
> fe80::20c:29ff:fe63:a673/128 :: U 0 0 1 lo
> fe80::/64 :: U 256 0 0 eth0
> ff00::/8 :: U 256 0 0 eth0

VE is missing default GW, I expect it to show there somehow automaticaly
like in HW node... maybe this is one of my bugs.

When I try to ping my HW node, for that I've routes there I get:

> root@ip6v2-test:/# ping6 2001:718:1:e::23:3004
> PING 2001:718:1:e::23:3004(2001:718:1:e::23:3004) 56 data bytes
> From 2001:718:1:e::23:3012 icmp_seq=1 Destination unreachable: Address unreachable
> From 2001:718:1:e::23:3012 icmp_seq=2 Destination unreachable: Address unreachable

and on HW node:

> staj-dev:~# tcpdump -i vzbr0
> tcpdump: WARNING: vzbr0: no IPv4 address assigned
> tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
> listening on vzbr0, link-type EN10MB (Ethernet), capture size 96 bytes
> 11:23:40.306129 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation, who has staj-dev6.cesnet.cz, length 32
> 11:23:41.305975 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation, who has staj-dev6.cesnet.cz, length 32

When I try oposite direction:

> staj-dev:~# ping6 2001:718:1:e::23:3012
> PING 2001:718:1:e::23:3012(2001:718:1:e::23:3012) 56 data bytes
> [... no output at all ...]

> root@ip6v2-test:/# tcpdump -n -i eth0
> tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
> listening on eth0, link-type EN10MB (Ethernet), capture size 96 bytes
> 09:32:01.755512 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation, who has 2001:718:1:e::23:3004, length 32
> 09:32:01.755721 IP6 2001:718:1:e::23:3004 > 2001:718:1:e::23:3012: ICMP6, echo request, seq 27, length 64
> 09:32:02.756372 IP6 2001:718:1:e::23:3004 > 2001:718:1:e::23:3012: ICMP6, echo request, seq 28, length 64
> 09:32:02.757356 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation, who has 2001:718:1:e::23:3004, length 32
> 09:32:03.756222 IP6 2001:718:1:e::23:3004 > 2001:718:1:e::23:3012: ICMP6, echo request, seq 29, length 64
> 09:32:03.757208 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation, who has 2001:718:1:e::23:3004, length 32
> 09:32:04.756070 IP6 2001:718:1:e::23:3004 > 2001:718:1:e::23:3012: ICMP6, echo request, seq 30, length 64
> 09:32:04.757057 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation, who has 2001:718:1:e::23:3004, length 32
> 09:32:05.755924 IP6 2001:718:1:e::23:3004 > 2001:718:1:e::23:3012: ICMP6, echo request, seq 31, length 64
> 09:32:06.755787 IP6 2001:718:1:e::23:3004 > 2001:718:1:e::23:3012: ICMP6, echo request, seq 32, length 64
> 09:32:06.756762 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation, who has 2001:718:1:e::23:3004, length 32

Any thoughts, please?
--

Jan Tomasek aka Semik
http://www.tomasek.cz/

Subject: Re: IPv6 again but this time with veth
Posted by gblond on Fri, 29 Jun 2007 08:43:32 GMT
View Forum Message <> Reply to Message

venet device works with IPv6!
Describe steps which you did for venet device.

On the Thursday 28 June 2007 22:44 Jan Tomasek, wrote:

> I need IPv4 and IPv6 support in VE.
>
>
> I would prefer venet interface, but that doesn't work. See thread:
> http://forum.openvz.org/index.php?t=msg&goto=12809&& amp;srch=IPv6#msg_12809

> Do you have any other suggestions?

--

Vitaliy Gusev,

Subject: Re: IPv6 again but this time with veth
Posted by gblond on Fri, 29 Jun 2007 08:50:22 GMT
View Forum Message <> Reply to Message

You added veth012.0 device to the bridge br0. How about eth0 ?

On the Friday 29 June 2007 12:23 Jan Tomasek, wrote:
> Oh. I got lost! Could I try to describe it again?
>
> My HW node is staj-dev:
>
> auto eth0
> iface eth0 inet static
> address 195.113.233.4
> netmask 255.255.255.0
> network 195.113.233.0
> broadcast 195.113.233.255
> gateway 195.113.233.1
> dns-nameservers 195.113.144.233 195.113.144.194
> dns-search cesnet.cz
>
> iface eth0 inet6 static
> post-up /sbin/sysctl -w "net.ipv6.conf.eth0.autoconf=0"
> address 2001:718:1:e::23:3004
> netmask 64
>
> My VE node, has this config (ip6v2-test):
>
> auto eth0
> iface eth0 inet static
> address 195.113.233.12
> netmask 255.255.255.0
> network 195.113.233.0
> broadcast 195.113.233.255
> gateway 195.113.233.4
> dns-nameservers 195.113.144.233 195.113.144.194
> dns-search cesnet.cz
>
>
> eth0 inet6 static

> iface eth0 inet6 static
> address 2001:718:1:e::23:3012
> netmask 64
>
> After executing:
>
> brctl addbr vzbr0
> brctl addif vzbr0 veth012.0
> ifconfig vzbr0 0
> echo 1 > /proc/sys/net/ipv4/conf/vzbr0/forwarding
> echo 1 > /proc/sys/net/ipv4/conf/vzbr0/proxy_arp
> echo 1 > /proc/sys/net/ipv4/conf/eth0/forwarding
> echo 1 > /proc/sys/net/ipv4/conf/eth0/proxy_arp
> ip route add 195.113.233.12 dev vzbr0
> ip route add 2001:718:1:e::23:3012 dev vzbr0
>
> **MARK**
>
> VE starts to be reachable from Internet on IPv4. What should I do next?
>
> echo 1 > /proc/sys/net/ipv6/conf/eth0/forwarding
>
> somehow broke IPv6 for HW node.
>
> echo 1 > /proc/sys/net/ipv6/conf/vzbr0/forwarding
>
> cause no problem.
>
> Well - returning to point marked as **MARK**.
>
> Route table on HW node:
>
> > staj-dev:~# route -6
> > Kernel IPv6 routing table
> > Destination Next Hop Flags Metric Ref Use Iface
> > ::1/128 :: U 0 6 1 lo
> > 2001:718:1:e::23:3004/128 :: U 0 788 1 lo
> > 2001:718:1:e::23:3012/128 :: U 1024 0 0 vzbr0
> > 2001:718:1:e::/64 :: U 256 169 0 eth0
> > fe80::/128 :: U 0 0 2 lo
> > fe80::1/128 :: U 0 0 1 lo
> > fe80::20c:29ff:fe63:a674/128 :: U 0 0 1 lo
> > fe80::20c:29ff:fe63:a674/128 :: U 0 0 1 lo
> > fe80::213:72ff:fe1b:b97/128 :: U 0 32 1 lo
> > fe80::/64 :: U 256 0 0 eth0
> > fe80::/64 :: U 256 0 0 venet0
> > fe80::/64 :: U 256 0 0 veth012.0
> > fe80::/64 :: U 256 0 0 vzbr0

> > f020::1/128 ff02::1 UC 0 165 0 eth0
> > ff00::/8 :: U 256 0 0 eth0
> > ff00::/8 :: U 256 0 0 venet0
> > ff00::/8 :: U 256 0 0 veth012.0
> > ff00::/8 :: U 256 0 0 vzbr0
> > ::/0 fe80::215:faff:fe87:3100 UGDA 1024 167 0 eth0
>
> Route table on VE node:
>
> > root@ip6v2-test:/# route -6
> > Kernel IPv6 routing table
> > Destination Next Hop Flags Metric Ref Use Iface
> > ::1/128 :: U 0 1 1 lo
> > 2001:718:1:e::23:3012/128 :: U 0 53 1 lo
> > 2001:718:1:e::/64 :: U 256 0 0 eth0
> > fe80::20c:29ff:fe63:a673/128 :: U 0 0 1 lo
> > fe80::/64 :: U 256 0 0 eth0
> > ff00::/8 :: U 256 0 0 eth0
>
> VE is missing default GW, I expect it to show there somehow automaticaly
> like in HW node... maybe this is one of my bugs.
>
> When I try to ping my HW node, for that I've routes there I get:
>
> > root@ip6v2-test:/# ping6 2001:718:1:e::23:3004
> > PING 2001:718:1:e::23:3004(2001:718:1:e::23:3004) 56 data bytes
> > From 2001:718:1:e::23:3012 icmp_seq=1 Destination unreachable: Address unreachable
> > From 2001:718:1:e::23:3012 icmp_seq=2 Destination unreachable: Address unreachable
>
> and on HW node:
>
> > staj-dev:~# tcpdump -i vzbr0
> > tcpdump: WARNING: vzbr0: no IPv4 address assigned
> > tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
> > listening on vzbr0, link-type EN10MB (Ethernet), capture size 96 bytes
> > 11:23:40.306129 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation, who has staj-dev6.cesnet.cz, length 32
> > 11:23:41.305975 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation, who has staj-dev6.cesnet.cz, length 32
>
> When I try oposite direction:
>
> > staj-dev:~# ping6 2001:718:1:e::23:3012
> > PING 2001:718:1:e::23:3012(2001:718:1:e::23:3012) 56 data bytes
> > [... no output at all ...]
>
> > root@ip6v2-test:/# tcpdump -n -i eth0
> > tcpdump: verbose output suppressed, use -v or -vv for full protocol decode

> > listening on eth0, link-type EN10MB (Ethernet), capture size 96 bytes
> > 09:32:01.755512 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation, who has 2001:718:1:e::23:3004, length 32
> > 09:32:01.755721 IP6 2001:718:1:e::23:3004 > 2001:718:1:e::23:3012: ICMP6, echo request, seq 27, length 64
> > 09:32:02.756372 IP6 2001:718:1:e::23:3004 > 2001:718:1:e::23:3012: ICMP6, echo request, seq 28, length 64
> > 09:32:02.757356 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation, who has 2001:718:1:e::23:3004, length 32
> > 09:32:03.756222 IP6 2001:718:1:e::23:3004 > 2001:718:1:e::23:3012: ICMP6, echo request, seq 29, length 64
> > 09:32:03.757208 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation, who has 2001:718:1:e::23:3004, length 32
> > 09:32:04.756070 IP6 2001:718:1:e::23:3004 > 2001:718:1:e::23:3012: ICMP6, echo request, seq 30, length 64
> > 09:32:04.757057 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation, who has 2001:718:1:e::23:3004, length 32
> > 09:32:05.755924 IP6 2001:718:1:e::23:3004 > 2001:718:1:e::23:3012: ICMP6, echo request, seq 31, length 64
> > 09:32:06.755787 IP6 2001:718:1:e::23:3004 > 2001:718:1:e::23:3012: ICMP6, echo request, seq 32, length 64
> > 09:32:06.756762 IP6 2001:718:1:e::23:3012 > ff02::1:ff23:3004: ICMP6, neighbor solicitation, who has 2001:718:1:e::23:3004, length 32
>
> Any thoughts, please?

--

Vitaliy Gusev,

Subject: Re: IPv6 again but this time with veth
Posted by Jan Tomasek on Fri, 29 Jun 2007 08:53:41 GMT
View Forum Message <> Reply to Message

Vitaliy Gusev wrote:
> venet device works with IPv6!
> Describe steps which you did for venet device.

I created VE 233005:
IPv4=195.113.233.5
IPv6=2001:718:1:e::23:3005

IPv4 is working ok.

When I try to add IPv6 I get:

staj-dev:~# vzctl set 233005 --ipadd '2001:718:1:e::23:3005' --save

Adding IP address(es): 2001:718:1:e::23:3005
Failed to bring up venet0:1.
Configure meminfo: 49152
Saved parameters for VE 233005

> staj-dev:~# vzctl enter 233005
> entered into VE 233005
> root@ipv6-pokus:/# ifconfig -a
> lo Link encap:Local Loopback
> inet addr:127.0.0.1 Mask:255.0.0.0
> inet6 addr: ::1/128 Scope:Host
> UP LOOPBACK RUNNING MTU:16436 Metric:1
> RX packets:0 errors:0 dropped:0 overruns:0 frame:0
> TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
> collisions:0 txqueuelen:0
> RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
>
> venet0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
> inet addr:127.0.0.1 P-t-P:127.0.0.1 Bcast:0.0.0.0 Mask:255.255.255.255
> UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
> RX packets:0 errors:0 dropped:0 overruns:0 frame:0
> TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
> collisions:0 txqueuelen:0
> RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
>
> venet0:0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
> inet addr:195.113.233.5 P-t-P:195.113.233.5 Bcast:0.0.0.0 Mask:255.255.255.255
> UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1

There is no venet0:1 inside of VE 233005.

--

Jan Tomasek aka Semik
http://www.tomasek.cz/

Subject: Re: IPv6 again but this time with veth
Posted by gblond on Fri, 29 Jun 2007 10:12:01 GMT
View Forum Message <> Reply to Message

Possible it is a bug in vzctl's configuration script. Which OS do you use for VE/Host? What is vzctl version?

On the Friday 29 June 2007 12:53 Jan Tomasek, wrote:
> Vitaliy Gusev wrote:
> > venet device works with IPv6!
> > Describe steps which you did for venet device.

>
> I created VE 233005:
> IPv4=195.113.233.5
> IPv6=2001:718:1:e::23:3005
>
> IPv4 is working ok.
>
> When I try to add IPv6 I get:
>
> staj-dev:~# vzctl set 233005 --ipadd '2001:718:1:e::23:3005' --save
> Adding IP address(es): 2001:718:1:e::23:3005
> Failed to bring up venet0:1.
> Configure meminfo: 49152
> Saved parameters for VE 233005
>
> > staj-dev:~# vzctl enter 233005
> > entered into VE 233005
> > root@ipv6-pokus:/# ifconfig -a
> > lo   Link encap:Local Loopback
> >   inet addr:127.0.0.1 Mask:255.0.0.0
> >   inet6 addr: ::1/128 Scope:Host
> >   UP LOOPBACK RUNNING MTU:16436 Metric:1
> >   RX packets:0 errors:0 dropped:0 overruns:0 frame:0
> >   TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
> >   collisions:0 txqueuelen:0
> >   RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
> >
> > venet0   Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
> >   inet addr:127.0.0.1 P-t-P:127.0.0.1 Bcast:0.0.0.0 Mask:255.255.255.255
> >   UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
> >   RX packets:0 errors:0 dropped:0 overruns:0 frame:0
> >   TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
> >   collisions:0 txqueuelen:0
> >   RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
> >
> > venet0:0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
> >   inet addr:195.113.233.5 P-t-P:195.113.233.5 Bcast:0.0.0.0 Mask:255.255.255.255
> >   UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
>
> There is no venet0:1 inside of VE 233005.
>

--

Vitaliy Gusev,

Subject: Re: IPv6 again but this time with veth
Posted by Jan Tomasek on Fri, 29 Jun 2007 10:15:51 GMT
View Forum Message <> Reply to Message

Vitaliy Gusev wrote:
> Possible it is a bug in vzctl's configuration script. Which OS do you
> use for VE/Host? What is vzctl version?

Both systems are Debian etch.

staj-dev:~# vzctl --version
vzctl version 3.0.16-5dso1

--

Jan Tomasek aka Semik
http://www.tomasek.cz/

Subject: Re: IPv6 again but this time with veth
Posted by Jan Tomasek on Fri, 29 Jun 2007 10:45:11 GMT
View Forum Message <> Reply to Message

Vitaliy Gusev wrote:
> You added veth012.0 device to the bridge br0. How about eth0 ?
>

It wasn't there. After adding it I finally could yell "IT WORSK!" but
just for while. :(In few minutes it stopped working on IPv4 (v6 is
still fine). I discovered that problem is in arp (my GW is .233.1):

> staj-dev:~# arp -n
> Address HWtype HWaddress Flags Mask Iface
> 195.113.233.224 (incomplete) eth0
> 195.113.233.12 ether 00:0C:29:63:A6:73 C vzbr0
> 195.113.233.254 (incomplete) eth0
> 195.113.233.92 (incomplete) eth0
> 195.113.233.193 (incomplete) eth0
> 195.113.233.99 ether 00:0C:29:6D:67:1A C vzbr0
> 195.113.233.99 (incomplete) eth0
> 195.113.233.36 (incomplete) eth0
> 195.113.233.254 ether 00:13:72:39:61:AA C vzbr0
> 195.113.233.2 (incomplete) eth0
> 195.113.233.1 (incomplete) eth0
> 195.113.233.245 (incomplete) eth0
> 195.113.233.246 (incomplete) eth0

staj-dev:~# arp -s 195.113.233.1 00:15:FA:87:31:0

> Address HWtype HWaddress Flags Mask Iface
> 195.113.233.91 (incomplete) eth0
> 195.113.233.246 (incomplete) eth0
> 195.113.233.34 (incomplete) eth0
> 195.113.233.12 ether 00:0C:29:63:A6:73 C vzbr0
> 195.113.233.74 (incomplete) eth0
> 195.113.233.1 ether 00:15:FA:87:31:00 CM eth0

fixes problem. But that is really ugly hack, this way I have to add
static arp records for every host in network 195.113.233.0/24 outside of
this HW node. That is not good way.

Any idea why ARP (and IPv4) stop work? Command I do execute are:

```
brctl addbr vzbr0
ifconfig vzbr0 0
echo 1 > /proc/sys/net/ipv4/conf/vzbr0/forwarding
echo 1 > /proc/sys/net/ipv4/conf/vzbr0/proxy_arp
brctl addif vzbr0 veth012.0
brctl addif vzbr0 eth0
ip route add 195.113.233.12 dev vzbr0
```

--

Jan Tomasek aka Semik
http://www.tomasek.cz/