Subject: IPv6 with venet
Posted by Jan Tomasek on Mon, 09 Jul 2007 13:40:18 GMT
View Forum Message <> Reply to Message

Hi,

after upgrade to vzctl version 3.0.18-1dso1 I got new set of
/etc/vz/dists/scripts/. There were changes in IPv6 code for Debian.

Thanks to those changes now IPv6 works between HW node and VE:

> staj-dev:~# ping6 2001:718:1:e::23:3005
> PING 2001:718:1:e::23:3005(2001:718:1:e::23:3005) 56 data bytes
> 64 bytes from 2001:718:1:e::23:3005: icmp_seq=1 ttl=64 time=0.068 ms

Problem is that I still can not reach VE from outside (staff is box with
IPv6 support).

> semik@staff:~$ ping6 2001:718:1:e::23:3005
> PING 2001:718:1:e::23:3005(2001:718:1:e::23:3005) 56 data bytes

ping does not show anything, it just hangs. In tcpdump on HW node I see:

> 15:38:41.465836 802.1d unknown version
> 15:38:41.498548 IP6 staff.cesnet.cz > 2001:718:1:e::23:3005: ICMP6, echo request, seq 923,
length 64
> 15:38:42.498349 IP6 staff.cesnet.cz > 2001:718:1:e::23:3005: ICMP6, echo request, seq 924,
length 64
> 15:38:43.462572 802.1d unknown version
> 15:38:43.498152 IP6 staff.cesnet.cz > 2001:718:1:e::23:3005: ICMP6, echo request, seq 925,
length 64
> 15:38:43.733136 arp who-has 195.113.233.73 tell 195.113.233.1
> 15:38:44.497951 IP6 staff.cesnet.cz > 2001:718:1:e::23:3005: ICMP6, echo request, seq 926,
length 64
> 15:38:45.427653 IP6 fe80::215:faff:fe87:3100 > ip6-allnodes: ICMP6, router advertisement,
length 64
> 15:38:45.462453 802.1d unknown version
> 15:38:45.497753 IP6 staff.cesnet.cz > 2001:718:1:e::23:3005: ICMP6, echo request, seq 927,
length 64
> 15:38:46.497552 IP6 staff.cesnet.cz > 2001:718:1:e::23:3005: ICMP6, echo request, seq 928,
length 64

I was suspecting route problems but routes does look ok. VE:

> root@ipv6-pokus:/# route -6
> Kernel IPv6 routing table
> Destination Next Hop Flags Metric Ref Use Iface
> ::1/128 :: U 0 1 1 lo

> ::1/128 :: U 0 0 1 lo
> 2001:718:1:e::23:3005/128 :: U 0 13 1 lo
> 2001:718:1:101:20b:dbff:fe93:5c38/128 2001:718:1:101:20b:dbff:fe93:5c38 UC 0
101 0 venet0
> fe80::/64 :: U 256 0 0 venet0
> ff00::/8 :: U 256 0 0 venet0
> ::/0 :: U 256 0 0 venet0

HW node:

> staj-dev:/proc/sys/net/ipv6/conf/venet0# route -6
> Kernel IPv6 routing table
> Destination Next Hop Flags Metric Ref Use Iface
> ::1/128 :: U 0 8 1 lo
> 2001:718:1:e::23:3004/128 :: U 0 926 1 lo
> 2001:718:1:e::23:3005/128 :: U 1024 0 0 venet0
> 2001:718:1:e:213:72ff:fe1b:b97/128 :: U 0 0 1 lo
> 2001:718:1:e::/64 :: U 256 302 0 eth0
> fe80::1/128 :: U 0 0 1 lo
> fe80::213:72ff:fe1b:b97/128 :: U 0 58 1 lo
> fe80::/64 :: U 256 0 0 eth0
> fe80::/64 :: U 256 0 0 venet0
> ff02::1/128 ff02::1 UC 0 327 0 eth0
> ff00::/8 :: U 256 0 0 eth0
> ff00::/8 :: U 256 0 0 venet0
> ::/0 fe80::215:faff:fe87:3100 UGDA 1024 483 0 eth0

Please, does anyone have any sugestions?

--

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

Subject: Re: IPv6 with venet
Posted by gblond on Mon, 09 Jul 2007 15:16:57 GMT
View Forum Message <> Reply to Message

Please check /proc/sys/net/ipv6/conf/all/forwarding

On the Monday 09 July 2007 17:40 Jan Tomasek, wrote:
> Hi,
>
> after upgrade to vzctl version 3.0.18-1dso1 I got new set of
> /etc/vz/dists/scripts/. There were changes in IPv6 code for Debian.
>
> Thanks to those changes now IPv6 works between HW node and VE:

>
> > staj-dev:~# ping6 2001:718:1:e::23:3005
> > PING 2001:718:1:e::23:3005(2001:718:1:e::23:3005) 56 data bytes
> > 64 bytes from 2001:718:1:e::23:3005: icmp_seq=1 ttl=64 time=0.068 ms
>
> Problem is that I still can not reach VE from outside (staff is box with
> IPv6 support).
>
> > semik@staff:~$ ping6 2001:718:1:e::23:3005
> > PING 2001:718:1:e::23:3005(2001:718:1:e::23:3005) 56 data bytes
>
> ping does not show anything, it just hangs. In tcpdump on HW node I see:
>
> > 15:38:41.465836 802.1d unknown version
> > 15:38:41.498548 IP6 staff.cesnet.cz > 2001:718:1:e::23:3005: ICMP6, echo request, seq 923,
length 64
> > 15:38:42.498349 IP6 staff.cesnet.cz > 2001:718:1:e::23:3005: ICMP6, echo request, seq 924,
length 64
> > 15:38:43.462572 802.1d unknown version
> > 15:38:43.498152 IP6 staff.cesnet.cz > 2001:718:1:e::23:3005: ICMP6, echo request, seq 925,
length 64
> > 15:38:43.733136 arp who-has 195.113.233.73 tell 195.113.233.1
> > 15:38:44.497951 IP6 staff.cesnet.cz > 2001:718:1:e::23:3005: ICMP6, echo request, seq 926,
length 64
> > 15:38:45.427653 IP6 fe80::215:faff:fe87:3100 > ip6-allnodes: ICMP6, router advertisement,
length 64
> > 15:38:45.462453 802.1d unknown version
> > 15:38:45.497753 IP6 staff.cesnet.cz > 2001:718:1:e::23:3005: ICMP6, echo request, seq 927,
length 64
> > 15:38:46.497552 IP6 staff.cesnet.cz > 2001:718:1:e::23:3005: ICMP6, echo request, seq 928,
length 64
>
> I was suspecting route problems but routes does look ok. VE:
>
> > root@ipv6-pokus:/# route -6
> > Kernel IPv6 routing table
> > Destination Next Hop Flags Metric Ref Use Iface
> > ::1/128 :: U 0 1 1 lo
> > ::1/128 :: U 0 0 1 lo
> > 2001:718:1:e::23:3005/128 :: U 0 13 1 lo
> > 2001:718:1:101:20b:dbff:fe93:5c38/128 2001:718:1:101:20b:dbff:fe93:5c38 UC 0
101 0 venet0
> > fe80::/64 :: U 256 0 0 venet0
> > ff00::/8 :: U 256 0 0 venet0
> > ::/0 :: U 256 0 0 venet0
>
> HW node:
>

> > staj-dev:/proc/sys/net/ipv6/conf/venet0# route -6
> > Kernel IPv6 routing table
> > Destination Next Hop Flags Metric Ref Use Iface
> > ::1/128 :: U 0 8 1 lo
> > 2001:718:1:e::23:3004/128 :: U 0 926 1 lo
> > 2001:718:1:e::23:3005/128 :: U 1024 0 0 venet0
> > 2001:718:1:e:213:72ff:fe1b:b97/128 :: U 0 0 1 lo
> > 2001:718:1:e::/64 :: U 256 302 0 eth0
> > fe80::1/128 :: U 0 0 1 lo
> > fe80::213:72ff:fe1b:b97/128 :: U 0 58 1 lo
> > fe80::/64 :: U 256 0 0 eth0
> > fe80::/64 :: U 256 0 0 venet0
> > ff02::1/128 ff02::1 UC 0 327 0 eth0
> > ff00::/8 :: U 256 0 0 eth0
> > ff00::/8 :: U 256 0 0 venet0
> > ::/0 fe80::215:faff:fe87:3100 UGDA 1024 483 0 eth0
>
> Please, does anyone have any sugestions?
>

--
Vitaliy Gusev

Subject: Re: IPv6 with venet
Posted by Jan Tomasek on Tue, 10 Jul 2007 07:40:11 GMT
View Forum Message <> Reply to Message

Vitaliy Gusev wrote:
> Please check /proc/sys/net/ipv6/conf/all/forwarding

It is disabled:

```
 staj-dev:~# for f in /proc/sys/net/ipv6/conf/*/forwarding; do echo $f;
cat $f; done
 /proc/sys/net/ipv6/conf/all/forwarding
 0
 /proc/sys/net/ipv6/conf/default/forwarding
 0
 /proc/sys/net/ipv6/conf/eth0/forwarding
 0
 /proc/sys/net/ipv6/conf/lo/forwarding
 0
 /proc/sys/net/ipv6/conf/venet0/forwarding
 0
```

Problem is that when I do:

```
echo "1" > /proc/sys/net/ipv6/conf/all/forwarding
```

HW node stops responding on IPv6 ping, reason probably is that setting
that removes default GW:

```
staj-dev:~# route -6
Kernel IPv6 routing table
Destination                    Next Hop
           Flags Metric Ref    Use Iface
::1/128                        ::
           U     0      13       1 lo
2001:718:1:e::/128                  ::
           U     0      0        2 lo
2001:718:1:e::23:3004/128           ::
           U     0      266      1 lo
2001:718:1:e::/64                   ::
           U     256    217      0 eth0
fe80::/128                          ::
           U     0      0        2 lo
fe80::1/128                         ::
           U     0      0        1 lo
fe80::213:72ff:fe1b:b97/128         ::
           U     0      8        1 lo
fe80::/64                           ::
           U     256    0        0 eth0
fe80::/64                           ::
           U     256    0        0 venet0
ff02::1/128                         ff02::1
           UC    0      229      0 eth0
ff00::/8                            ::
           U     256    0        0 eth0
ff00::/8                            ::
           U     256    0        0 venet0
```

When I set GW manualy it start work and it works also for VE. GREAT!

Well and now. Why setting forwarding=1 removes GW? Is there any way how
to add it there automaticaly without my attention? I'm afraid that GW IP
fe80::215:faff:fe87:3100 is somehow related to router MAC and that it
will change when router become replaced. Any sugestion on this?

Thanks a lot!
--

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

Subject: Re: IPv6 with venet
Posted by Thorsten Schifferdeck on Tue, 10 Jul 2007 10:13:21 GMT
View Forum Message <> Reply to Message

Hi Jan,

Jan Tomasek schrieb:
> Vitaliy Gusev wrote:
>> Please check /proc/sys/net/ipv6/conf/all/forwarding
> [...]
> Problem is that when I do:
>
> echo "1" > /proc/sys/net/ipv6/conf/all/forwarding
>
> HW node stops responding on IPv6 ping, reason probably is that setting
> that removes default GW:

"Router Advertisements" (=autoconfiguration) are ignored on setting
forward to 1 to IPv6 proc settings and your (automatic) default route (
::0), is set to a RA IPv6 (link-local) address, is gone.

for Debian you can set IPv6 (addy and route) on your VE0 (hardware node):

--8<--(/etc/network/interfaces)--
[...]
iface <DEVICE>
[...]
iface <DEVICE> inet6 static
 address <IPv6_ADDRESS>
 netmask <NETMASK>
 gateway <IPv6_ADDRESS_of_your_GATEWAY>
 # set IPv6 forwarding here or in /etc/sysctl.conf
 # read man page of interfaces (5)
[...]
-->8--

Regards,
Thorsten Schifferdecker