
Subject: 6to4 - tunnel for VE

Posted by [shion](#) on Tue, 11 Aug 2009 15:07:06 GMT

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

Hello,

I need a separate 6to4 tunnel IP-address for a VE.

The HN and each VE have a own static public IPv4-address.

Configuration on HN:

```
auto sit1
```

```
iface sit1 inet6 v4tunnel
```

```
    address 2002:75aa:2a14::1
```

```
    netmask 64
```

```
    gateway ::192.88.99.1
```

```
    endpoint any
```

```
    local 117.170.42.20
```

Enable IPv6 in vz.conf

```
ipv6="yes"
```

Enable IPv6 forwarding

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

After this configuration is a ping6 from 2002:75aa:2a14::1 possible.

```
# ping6 ipv6.google.com
```

```
PING ipv6.google.com(ey-in-x68.google.com) 56 data bytes
```

```
64 bytes from ey-in-x68.google.com: icmp_seq=1 ttl=54 time=42.0 ms
```

```
64 bytes from ey-in-x68.google.com: icmp_seq=2 ttl=54 time=41.6 ms
```

Next I have executed

```
vzctl set 111 --ipadd 2002:75aa:2a14::2 --save so that I can ping6 from the VE too.
```

The VE uses the venet0 interface.

So far, so good..

That works for now but I need another solution.

The VE has an own public IPv4-address and this address should be used for VEs 6to4 tunnel.

I tried the following things:

Set NEIGHBOUR_DEVS to all in vz.conf

```
NEIGHBOUR_DEVS=all
```

Set the desired 6to4 IPv6 address and removed the old IP:

```
vzctl set 111 --ipdel 2002:75aa:2a14::2 -save
```

```
vzctl set 111 --ipadd 2002:75aa:2a15::1 -save
```

Results:

- ping6 from HN works furthermore

- ping6 from VE doesn't work anymore

```
# ping6 ipv6.google.com
PING ipv6.google.com(ey-in-x68.google.com) 56 data bytes
^C
# tcpdump -n icmp6
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
16:24:39.462282 IP6 2002:75aa:2a15::1 > 2001:4860:a005::68: ICMP6, echo request, seq 1,
length 64
16:24:40.475794 IP6 2002:75aa:2a15::1 > 2001:4860:a005::68: ICMP6, echo request, seq 2,
length 64
16:24:41.475795 IP6 2002:75aa:2a15::1 > 2001:4860:a005::68: ICMP6, echo request, seq 3,
length 64
The ICMP reply packets are missing. So it could be a routing issue?!
```

VE - route -6

```
# route -6
```

Kernel IPv6 routing table

Destination	Next Hop	Flag Met Ref Use If
::1/128	::	U 256 0 0 venet0
fe80::/64	::	U 256 0 0 venet0
::/0	::	U 256 0 1 venet0
::/0	::	!n -1 1 1345 lo
::1/128	::	Un 0 1 425 lo
::1/128	::	Un 0 1 0 lo
2002:75aa:2a15::1/128	::	Un 0 1 4642 lo
ff00::/8	::	U 256 0 0 venet0
::/0	::	!n -1 1 1345 lo

HN - route -6

```
# route -6
```

Kernel IPv6 routing table

Destination	Next Hop	Flag Met Ref Use If
::192.88.99.1/128	::	U 1024 0 3 sit1
::/96	::	Un 256 0 0 sit1
2002:75aa:2a14::/64	::	U 256 0 0 sit1
2002:75aa:2a15::1/128	::	U 1024 0 0 venet0
fe80::1/128	::	U 256 0 0 venet0
fe80::/64	::	U 256 0 0 eth0
fe80::/64	::	U 256 0 0 venet0
fe80::/64	::	U 256 0 0 sit1
::/0	::192.88.99.1	UG 1024 0 7 sit1
::/0	::	!n -1 1301588 lo
::1/128	::	Un 0 1 31413 lo
::117.170.42.20/128	::	Un 0 1 0 lo
2002:75aa:2a14::/128	::	Un 0 1 0 lo
2002:75aa:2a14::1/128	::	Un 0 1 0 lo
fe80::/128	::	Un 0 1 0 lo
fe80::/128	::	Un 0 1 0 lo
fe80::1/128	::	Un 0 1 0 lo

```
fe80::214:22ff:fe73:78e4/128 ::          Un  0  1  0 lo
fe80::218:51ff:fec5:4e68/128 ::          Un  0  1  0 lo
ff00::/8                ::              U   256 0   0 eth0
ff00::/8                ::              U   256 0   0 venet0
ff00::/8                ::              U   256 0   0 sit1
::/0                    ::              !n  -1 1301588 lo
```

But it is possible to ping6 between HN <--> VE.

```
# ping6 2002:75aa:2a14::1
```

```
PING 2002:d946:8e2a::1(2002:d946:8e2a::1) 56 data bytes
```

```
64 bytes from 2002:75aa:2a14::1: icmp_seq=1 ttl=64 time=0.074 ms
```

```
64 bytes from 2002:75aa:2a14::1: icmp_seq=2 ttl=64 time=0.046 ms# ping6 2002:75aa:2a15::1
```

```
PING 2002:d946:8e51::1(2002:d946:8e51::1) 56 data bytes
```

```
64 bytes from 2002:75aa:2a15::1: icmp_seq=1 ttl=64 time=0.071 ms
```

```
64 bytes from 2002:75aa:2a15::1: icmp_seq=2 ttl=64 time=0.065 ms
```

versions:

```
# uname -r
```

```
2.6.26-2-openvz-amd64
```

```
# vzctl --version
```

```
vzctl version 3.0.22
```

So where is the problem?

I hope you can understand my problem.

If not feel free to ask, if something isn't clear.

Best regards

shion

(all IP-addresses are replaced)

Subject: Re: 6to4 - tunnel for VE

Posted by [maratrus](#) on Thu, 13 Aug 2009 11:56:19 GMT

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

Hello,

where tcpdump output was taken from?

Could you please add the following rule on the HN:

```
# ip neigh add proxy 2002:75aa:2a15::1 dev eth0
```

Do you have any firewall rules installed?

Could you please take tcpdump output from venet0 interfaces from inside the VE and from the HN (please, catch all proto not only icmp).

Subject: Re: 6to4 - tunnel for VE
Posted by [shion](#) on Thu, 20 Aug 2009 14:37:59 GMT
[View Forum Message](#) <> [Reply to Message](#)

The tcpdump was taken from the VE and there are no firewall rules installed on HN and VE.

Well, I have added your rule on the HN but I'm not sure if it is active, because I get only a empty result with the following command.

```
# ip -6 neigh show
```

I have found this comment:

"Strangely, the ip neighbor show command does not display any entries added and deleted with ip neighbor add proxy, so arp is required to view these entries. In short, don't use ip neighbor add proxy. "

"#arp -a -n" shows me only IPv4 addresses. Either "ip neighbor add proxy" wasn't successful or I do something wrong with arp.

Is there another way to control whether the command "ip neigh add proxy" was successful?

Ping from VE

```
# ping6 ipv6.google.com
```

```
PING ipv6.google.com(ey-in-x68.google.com) 56 data bytes
```

```
^C
```

```
--- ipv6.google.com ping statistics ---
```

```
6 packets transmitted, 0 received, 100% packet loss, time 5030ms
```

tcpdump from VE

```
# tcpdump -i venet0 -n
```

```
tcpdump: WARNING: arptype 65535 not supported by libpcap - falling back to cooked socket
```

```
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
```

```
13:33:29.115424 IP 117.170.42.21.41339 > 117.170.42.99.53: 34024+ AAAA? ipv6.google.com.  
(33)
```

```
13:33:29.115653 IP 117.170.42.99.53 > 117.170.42.21.41339: 34024 2/0/0 CNAME[domain]
```

```
13:33:29.116173 IP 117.170.42.21.37874 > 117.170.42.99.53: 6854+[domain]
```

```
13:33:29.116554 IP 117.170.42.99.53 > 117.170.42.21.37874: 6854[domain]
```

```
13:33:29.117074 IP6 2002:75aa:2a15::1 > 2001:4860:a005::68: ICMP6, echo request, seq 1,  
length 64
```

```
13:33:29.159630 IP 192.88.99.1 > 117.170.42.21: IP6 2001:4860:a005::68 > 2002:75aa:2a15::1:  
ICMP6, echo reply, seq 1, length 64
```

```
13:33:29.159663 IP 117.170.42.21 > 192.88.99.1: ICMP 117.170.42.21 protocol 41 port 0  
unreachable, length 132
```

```
13:33:30.131819 IP6 2002:75aa:2a15::1 > 2001:4860:a005::68: ICMP6, echo request, seq 2,  
length 64
```

```
13:33:30.174764 IP 192.88.99.1 > 117.170.42.21: IP6 2001:4860:a005::68 > 2002:75aa:2a15::1:  
ICMP6, echo reply, seq 2, length 64
```

```
13:33:30.174808 IP 117.170.42.21 > 192.88.99.1: ICMP 117.170.42.21 protocol 41 port 0  
unreachable, length 132
```

```
13:33:31.135296 IP6 2002:75aa:2a15::1 > 2001:4860:a005::68: ICMP6, echo request, seq 3,  
length 64
```

```
13:33:31.178152 IP 192.88.99.1 > 117.170.42.21: IP6 2001:4860:a005::68 > 2002:75aa:2a15::1:
ICMP6, echo reply, seq 3, length 64
13:33:31.178189 IP 117.170.42.21 > 192.88.99.1: ICMP 117.170.42.21 protocol 41 port 0
unreachable, length 132
13:33:32.135325 IP6 2002:75aa:2a15::1 > 2001:4860:a005::68: ICMP6, echo request, seq 4,
length 64
[...]
```

tcpdump from HN

```
# tcpdump -i venet0 -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
```

```
15:33:29.115424 IP 117.170.42.21.41339 > 117.170.42.99.53: 34024+ AAAA? ipv6.google.com.
(33)
```

```
15:33:29.115653 IP 117.170.42.99.53 > 117.170.42.21.41339: 34024 2/0/0 CNAME[domain]
```

```
15:33:29.116173 IP 117.170.42.21.37874 > 117.170.42.99.53: 6854+[domain]
```

```
15:33:29.116554 IP 117.170.42.99.53 > 117.170.42.21.37874: 6854+[domain]
```

```
15:33:29.119851 IP6 2002:75aa:2a15::1 > 2001:4860:a005::68: ICMP6, echo request, seq 1,
length 64
```

```
15:33:29.159630 IP 192.88.99.1 > 117.170.42.21: IP6 2001:4860:a005::68 > 2002:75aa:2a15::1:
ICMP6, echo reply, seq 1, length 64
```

```
15:33:29.159663 IP 117.170.42.21 > 192.88.99.1: ICMP 117.170.42.21 protocol 41 port 0
unreachable, length 132
```

```
15:33:30.161829 IP6 2002:75aa:2a15::1 > 2001:4860:a005::68: ICMP6, echo request, seq 2,
length 64
```

```
15:33:30.174764 IP 192.88.99.1 > 117.170.42.21: IP6 2001:4860:a005::68 > 2002:75aa:2a15::1:
ICMP6, echo reply, seq 2, length 64
```

```
15:33:30.174808 IP 117.170.42.21 > 192.88.99.1: ICMP 117.170.42.21 protocol 41 port 0
unreachable, length 132
```

```
15:33:31.176834 IP6 2002:75aa:2a15::1 > 2001:4860:a005::68: ICMP6, echo request, seq 3,
length 64
```

```
15:33:31.178152 IP 192.88.99.1 > 117.170.42.21: IP6 2001:4860:a005::68 > 2002:75aa:2a15::1:
ICMP6, echo reply, seq 3, length 64
```

```
15:33:31.178189 IP 117.170.42.21 > 192.88.99.1: ICMP 117.170.42.21 protocol 41 port 0
unreachable, length 132
```

```
[...]
```

If I remember correctly, 6to4 tunneling is also referred as proto-41.

But I don't know the message "protocol 41 port 0 unreachable" and I don't find any helpful information about it with google.