
Subject: IPv6 ??

Posted by [Benoit Branciard](#) on Thu, 24 Jan 2008 17:29:03 GMT

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

Temporarily giving up messing with veth devices (forgot dhcpd server for now), I concentrated on setting up a basic VS with both IPv4 and IPv6 addresses, using standard venet device.

I'm using Debian Etch hardware node with debian.systs.org packages.

So I did:

- install packages and reboot as necessary :

```
fzakernel-2.6.18-686 vzctl vzctl-ostmpl-debian vdump vzprocps vzquota
```

- ensure `/etc/sysctl.conf` contains :

```
net.ipv4.conf.default.forwarding = 1
net.ipv4.conf.all.rp_filter = 1
kernel.sysrq = 1
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0
```

- ensure `/etc/vz/vz.conf` contains:

```
IPV6="yes"
```

- create VS :

```
vzctl create 112 --ipadd XXX.YYY.ZZZ.12 --ipadd xxx:yyy:zzz::12
--hostname brahma-vps112
```

(default template is `debian-4.0-i386-minimal`, hardware node has only one physical interface activated, on the same subnet as VS)

- start VS :

```
vzctl start 112
```

Result:

- IPv4 VS works as expected (ping succeeds from hardware node, external server in same subnet, external server on another subnet).

- IPv6 VS doesn't work !!! (ping6 works **only** between VS and hardware node, not from external sources; VS has IPv6 address correctly defined, as per `ifconfig`).

Am I missing something ?

I suspected it may be useful to set `sysctl net.ipv6.conf.<iface>.forwarding` to 1, but :

- this turns the box in "router" mode, disabling RA acceptance from default router (annoying)
- this doesn't work either !!

--

Ce message a ete verifie par MailScanner pour des virus ou des polluriels et rien de suspect n'a ete trouve.

Subject: Re: IPv6 ??

Posted by [Jan Tomasek](#) on Thu, 24 Jan 2008 20:47:44 GMT

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

Hello Benoit,

> Temporarily giving up messing with veth devices (forgot dhcpd server for
> now), I concentrated on setting up a basic VS with both IPv4 and IPv6
> addresses, using standard venet device.

>

> [...]

> Result:

>

> - IPv4 VS works as expected (ping succeeds from hardware node, external
> server in same subnet, external server on another subnet).

>

> - IPv6 VS doesn't work !!! (ping6 works *only* between VS and hardware
> node, not from external sources; VS has IPv6 address correctly defined,
> as per ifconfig).

>

> Am I missing something ?

>

> I suspected it may be useful to set `sysctl`

> `net.ipv6.conf.<iface>.forwarding` to 1, but :

> - this turns the box in "router" mode, disabling RA acceptance from
> default router (annoying)

> - this doesn't work either !!

Very interesting! I'm having exactly same problem here. IPv6 was working perfectly for me after some problems before several months. Now I had to reboot my HW node and IPv6 is gone.

When I try to ping VE from real world I'm getting:

```
> semik@staff:~$ ping6 2001:718:1:e::23:3257
> PING 2001:718:1:e::23:3257(2001:718:1:e::23:3257) 56 data bytes
> From 2001:718:1:101::1 icmp_seq=0 Destination unreachable: Address unreachable
> From 2001:718:1:101::1 icmp_seq=3 Destination unreachable: Address unreachable
```

From HW node to VE or from one VE to another VE is IPv6 working fine. I think that HW node didn't sent some info to it's IPv6 gateway about presence of new IPv6 address. Sadly I'm clue less how to better debug this. But!

I've discovered trick. Run alternate version of this on your HW node:

```
> ifconfig eth0 add 2001:718:1:e::23:3257/64
> sleep 5
> ping6 -I 2001:718:1:e::23:3257 -c 1 2001:718:1:e::1
> ifconfig eth0 del 2001:718:1:e::23:3257/64
```

2001:718:1:e::23:3257 - IPv6 address of VE
2001:718:1:e::1 - gateway

That will cause that router will learn about presence of VE. It lasts for several hours. I discovered it at about 17 clock and now at 22 are VE still IPv6 online... but I'm afraid that clearing IPv6 neighbours table on router will kill this. Something is not working on HW node.

I'm running 2.6.18-028stab051 on Debian Etch. Output of `sysctl -a |grep net.ipv6 | sort >/tmp/chlivek.ipv6` is attached. Hope someone will point out where problem is. Yesterday I've reboot my second HW node...

--

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

```
net.ipv6.bindv6only = 0
net.ipv6.conf.all.accept_ra = 1
net.ipv6.conf.all.accept_ra_defrtr = 1
net.ipv6.conf.all.accept_ra_pinfo = 1
net.ipv6.conf.all.accept_ra_rtr_pref = 1
net.ipv6.conf.all.accept_redirects = 1
net.ipv6.conf.all.autoconf = 1
net.ipv6.conf.all.dad_transmits = 1
net.ipv6.conf.all.force_mld_version = 0
net.ipv6.conf.all.forwarding = 1
net.ipv6.conf.all.hop_limit = 64
net.ipv6.conf.all.max_addresses = 16
net.ipv6.conf.all.max_desync_factor = 600
```

```
net.ipv6.conf.all.mtu = 1280
net.ipv6.conf.all.regen_max_retry = 5
net.ipv6.conf.all.router_probe_interval = 60
net.ipv6.conf.all.router_solicitation_delay = 1
net.ipv6.conf.all.router_solicitation_interval = 4
net.ipv6.conf.all.router_solicitations = 3
net.ipv6.conf.all.temp_preferred_lft = 86400
net.ipv6.conf.all.temp_valid_lft = 604800
net.ipv6.conf.all.use_tempaddr = 0
net.ipv6.conf.default.accept_ra = 1
net.ipv6.conf.default.accept_ra_defrtr = 1
net.ipv6.conf.default.accept_ra_pinfo = 1
net.ipv6.conf.default.accept_ra_rtr_pref = 1
net.ipv6.conf.default.accept_redirects = 1
net.ipv6.conf.default.autoconf = 1
net.ipv6.conf.default.dad_transmits = 1
net.ipv6.conf.default.force_mld_version = 0
net.ipv6.conf.default.forwarding = 1
net.ipv6.conf.default.hop_limit = 64
net.ipv6.conf.default.max_addresses = 16
net.ipv6.conf.default.max_desync_factor = 600
net.ipv6.conf.default.mtu = 1280
net.ipv6.conf.default.regen_max_retry = 5
net.ipv6.conf.default.router_probe_interval = 60
net.ipv6.conf.default.router_solicitation_delay = 1
net.ipv6.conf.default.router_solicitation_interval = 4
net.ipv6.conf.default.router_solicitations = 3
net.ipv6.conf.default.temp_preferred_lft = 86400
net.ipv6.conf.default.temp_valid_lft = 604800
net.ipv6.conf.default.use_tempaddr = 0
net.ipv6.conf.eth0.accept_ra = 1
net.ipv6.conf.eth0.accept_ra_defrtr = 1
net.ipv6.conf.eth0.accept_ra_pinfo = 1
net.ipv6.conf.eth0.accept_ra_rtr_pref = 1
net.ipv6.conf.eth0.accept_redirects = 1
net.ipv6.conf.eth0.autoconf = 0
net.ipv6.conf.eth0.dad_transmits = 1
net.ipv6.conf.eth0.force_mld_version = 0
net.ipv6.conf.eth0.forwarding = 1
net.ipv6.conf.eth0.hop_limit = 64
net.ipv6.conf.eth0.max_addresses = 16
net.ipv6.conf.eth0.max_desync_factor = 600
net.ipv6.conf.eth0.mtu = 1500
net.ipv6.conf.eth0.regen_max_retry = 5
net.ipv6.conf.eth0.router_probe_interval = 60
net.ipv6.conf.eth0.router_solicitation_delay = 1
net.ipv6.conf.eth0.router_solicitation_interval = 4
net.ipv6.conf.eth0.router_solicitations = 3
```

```
net.ipv6.conf.eth0.temp_prefered_lft = 86400
net.ipv6.conf.eth0.temp_valid_lft = 604800
net.ipv6.conf.eth0.use_tempaddr = 0
net.ipv6.conf.lo.accept_ra = 1
net.ipv6.conf.lo.accept_ra_defrtr = 1
net.ipv6.conf.lo.accept_ra_pinfo = 1
net.ipv6.conf.lo.accept_ra_rtr_pref = 1
net.ipv6.conf.lo.accept_redirects = 1
net.ipv6.conf.lo.autoconf = 1
net.ipv6.conf.lo.dad_transmits = 1
net.ipv6.conf.lo.force_mld_version = 0
net.ipv6.conf.lo.forwarding = 1
net.ipv6.conf.lo.hop_limit = 64
net.ipv6.conf.lo.max_addresses = 16
net.ipv6.conf.lo.max_desync_factor = 600
net.ipv6.conf.lo.mtu = 16436
net.ipv6.conf.lo.regen_max_retry = 5
net.ipv6.conf.lo.router_probe_interval = 60
net.ipv6.conf.lo.router_solicitation_delay = 1
net.ipv6.conf.lo.router_solicitation_interval = 4
net.ipv6.conf.lo.router_solicitations = 3
net.ipv6.conf.lo.temp_prefered_lft = 86400
net.ipv6.conf.lo.temp_valid_lft = 604800
net.ipv6.conf.lo.use_tempaddr = -1
net.ipv6.conf.venet0.accept_ra = 1
net.ipv6.conf.venet0.accept_ra_defrtr = 1
net.ipv6.conf.venet0.accept_ra_pinfo = 1
net.ipv6.conf.venet0.accept_ra_rtr_pref = 1
net.ipv6.conf.venet0.accept_redirects = 1
net.ipv6.conf.venet0.autoconf = 1
net.ipv6.conf.venet0.dad_transmits = 1
net.ipv6.conf.venet0.force_mld_version = 0
net.ipv6.conf.venet0.forwarding = 1
net.ipv6.conf.venet0.hop_limit = 64
net.ipv6.conf.venet0.max_addresses = 16
net.ipv6.conf.venet0.max_desync_factor = 600
net.ipv6.conf.venet0.mtu = 1500
net.ipv6.conf.venet0.regen_max_retry = 5
net.ipv6.conf.venet0.router_probe_interval = 60
net.ipv6.conf.venet0.router_solicitation_delay = 1
net.ipv6.conf.venet0.router_solicitation_interval = 4
net.ipv6.conf.venet0.router_solicitations = 3
net.ipv6.conf.venet0.temp_prefered_lft = 86400
net.ipv6.conf.venet0.temp_valid_lft = 604800
net.ipv6.conf.venet0.use_tempaddr = 0
net.ipv6.icmp.ratelimit = 1000
net.ipv6.ip6frag_high_thresh = 262144
net.ipv6.ip6frag_low_thresh = 196608
```

```
net.ipv6.ip6frag_secret_interval = 600
net.ipv6.ip6frag_time = 60
net.ipv6.mld_max_msf = 64
net.ipv6.neigh.default.anycast_delay = 99
net.ipv6.neigh.default.app_solicit = 0
net.ipv6.neigh.default.base_reachable_time = 30
net.ipv6.neigh.default.base_reachable_time_ms = 30000
net.ipv6.neigh.default.delay_first_probe_time = 5
net.ipv6.neigh.default.gc_interval = 30
net.ipv6.neigh.default.gc_stale_time = 60
net.ipv6.neigh.default.gc_thresh1 = 128
net.ipv6.neigh.default.gc_thresh2 = 512
net.ipv6.neigh.default.gc_thresh3 = 1024
net.ipv6.neigh.default.locktime = 0
net.ipv6.neigh.default.mcast_solicit = 3
net.ipv6.neigh.default.proxy_delay = 79
net.ipv6.neigh.default.proxy_qlen = 64
net.ipv6.neigh.default.retrans_time = 1000
net.ipv6.neigh.default.retrans_time_ms = 1000
net.ipv6.neigh.default.unicast_solicit = 3
net.ipv6.neigh.default.unres_qlen = 3
net.ipv6.neigh.eth0.anycast_delay = 99
net.ipv6.neigh.eth0.app_solicit = 0
net.ipv6.neigh.eth0.base_reachable_time = 30
net.ipv6.neigh.eth0.base_reachable_time_ms = 30000
net.ipv6.neigh.eth0.delay_first_probe_time = 5
net.ipv6.neigh.eth0.gc_stale_time = 60
net.ipv6.neigh.eth0.locktime = 0
net.ipv6.neigh.eth0.mcast_solicit = 3
net.ipv6.neigh.eth0.proxy_delay = 79
net.ipv6.neigh.eth0.proxy_qlen = 64
net.ipv6.neigh.eth0.retrans_time = 1000
net.ipv6.neigh.eth0.retrans_time_ms = 1000
net.ipv6.neigh.eth0.unicast_solicit = 3
net.ipv6.neigh.eth0.unres_qlen = 3
net.ipv6.neigh.lo.anycast_delay = 99
net.ipv6.neigh.lo.app_solicit = 0
net.ipv6.neigh.lo.base_reachable_time = 30
net.ipv6.neigh.lo.base_reachable_time_ms = 30000
net.ipv6.neigh.lo.delay_first_probe_time = 5
net.ipv6.neigh.lo.gc_stale_time = 60
net.ipv6.neigh.lo.locktime = 0
net.ipv6.neigh.lo.mcast_solicit = 3
net.ipv6.neigh.lo.proxy_delay = 79
net.ipv6.neigh.lo.proxy_qlen = 64
net.ipv6.neigh.lo.retrans_time = 1000
net.ipv6.neigh.lo.retrans_time_ms = 1000
net.ipv6.neigh.lo.unicast_solicit = 3
```

```
net.ipv6.neigh.lo.unres_qlen = 3
net.ipv6.neigh.venet0.anycast_delay = 99
net.ipv6.neigh.venet0.app_solicit = 0
net.ipv6.neigh.venet0.base_reachable_time = 30
net.ipv6.neigh.venet0.base_reachable_time_ms = 30000
net.ipv6.neigh.venet0.delay_first_probe_time = 5
net.ipv6.neigh.venet0.gc_stale_time = 60
net.ipv6.neigh.venet0.locktime = 0
net.ipv6.neigh.venet0.mcast_solicit = 3
net.ipv6.neigh.venet0.proxy_delay = 79
net.ipv6.neigh.venet0.proxy_qlen = 64
net.ipv6.neigh.venet0.retrans_time = 1000
net.ipv6.neigh.venet0.retrans_time_ms = 1000
net.ipv6.neigh.venet0.ucast_solicit = 3
net.ipv6.neigh.venet0.unres_qlen = 3
net.ipv6.route.gc_elasticity = 0
net.ipv6.route.gc_interval = 30
net.ipv6.route.gc_min_interval = 0
net.ipv6.route.gc_min_interval_ms = 500
net.ipv6.route.gc_thresh = 1024
net.ipv6.route.gc_timeout = 60
net.ipv6.route.max_size = 4096
net.ipv6.route.min_adv_mss = 1
net.ipv6.route.mtu_expires = 600
```

Subject: Re: IPv6 ??

Posted by [Thorsten Schifferdeck](#) on Thu, 24 Jan 2008 20:57:28 GMT

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

Hi,

for official IPv6 i'm using sixxs as tunnelbroker,
these steps i've done connect a Debian OpenVZ system
to a IPv6 Network for the VE with a venet0 interface:

1) Setup an IPv6 Tunnel (-> <https://noc.sixxs.net/>)

Relevant IPv6 settings in "/etc/network/interfaces"

[...]

```
# device: sixxs ipv6 tunnel with static IPv6
```

```
auto sixxs
```

```
iface sixxs inet6 v4tunnel
```

```
    address <your_assigned_IPv6_ADDRESS>
```

```
    netmask 64
```

```
    endpoint <IPv4_ADDRESS_of_tunnelbroker>
```

```
    ttl 64
```

```
    up ip link set mtu 1280 dev sixxs
```

```
up ip route add default via <IPv6_ADDRESS_of_tunnelbroker> dev sixxs  
[...]
```

```
VE_0$ ifup sixxs
```

2) Setup OpenVZ with IPv6 :

```
# enable IPv6 forwarding for all Interfaces
```

```
VE_0$ sysctl -q -w net.ipv6.conf.all.forwarding=1
```

```
# enable IPv6 in /etc/vz/vz.conf -> IPV6="yes"
```

```
# and restart OpenVZ Subsystem, so venet has a IPv6 Address
```

```
VE_0$ /etc/init.d/vz restart
```

```
# give a VE an IPv6 address, you need a IPv6 subnet (
```

```
https://noc.sixxs.net/ )
```

```
VE_0$ vzctl set <VEID> --ipadd <IPv6_ADDRESS> [--yes]
```

3) Test it

If you connected to an official IPv6 network, you can test it, by pinging debian.systs.org (2001:6f8:109a::1a01) or visit IPv6 sites like <http://www.kame.net> and see the dancing turtle ;-)

Bye,

Thorsten

--

Thorsten Schifferdecker

tsd@debian.systs.org

---//---

Am Do, 24.01.2008, 18:29, schrieb Benoit Branciard:

```
> Temporarily giving up messing with veth devices (forgot dhcpd server for  
> now), I concentrated on setting up a basic VS with both IPv4 and IPv6  
> addresses, using standard venet device.
```

```
>
```

```
> I'm using Debian Etch hardware node with debian.systs.org packages.
```

```
>
```

```
> So I did:
```

```
>
```

```
> - install packages and reboot as necessary :
```

```
> fzakernel-2.6.18-686 vzctl vzctl-ostmpl-debian vzdump vzprocps vzquota
```

```
>
```

```
> - ensure /etc/sysctl.conf contains :
```

```
>
```

```
> net.ipv4.conf.default.forwarding = 1
```

```
> net.ipv4.conf.all.rp_filter = 1
```

```
> kernel.sysrq = 1
```



```
> net.ipv4.conf.default.send_redirects = 1
> net.ipv4.conf.all.send_redirects = 0
>
> - ensure /etc/vz/vz.conf contains:
> IPV6="yes"
>
> - create VS :
> vzctl create 112 --ipadd XXX.YYY.ZZZ.12 --ipadd xxx:yyy:zzzz::12
> --hostname brahma-vps112
>
> (default template is debian-4.0-i386-minimal, hardware node has only one
> physical interface activated, on the same subnet as VS)
>
> - start VS :
> vzctl start 112
>
>
> Result:
>
> - IPv4 VS works as expected (ping succeeds from hardware node, external
> server in same subnet, external server on another subnet).
>
> - IPv6 VS doesn't work !!! (ping6 works *only* between VS and hardware
> node, not from external sources; VS has IPv6 address correctly defined,
> as per ifconfig).
>
> Am I missing something ?
>
> I suspected it may be useful to set sysctl
> net.ipv6.conf.<iface>.forwarding to 1, but :
> - this turns the box in "router" mode, disabling RA acceptance from
> default router (annoying)
> - this doesn't work either !!
```

Subject: Re: IPv6 ??

Posted by [Thorsten Schifferdeck\[1\]](#) on Fri, 25 Jan 2008 06:30:42 GMT

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

Hi,

@jan: how get your IPv6 connectivity ?

IPv6 traceroute from noc.sixxs.net @ SixXS NOC, AS12871 to 2001:718:1:e::23:3257 :

Hop	Node	Loss%	Sent	Last	Avg	Best	Worst
StDev	ASN	Organisation					

```

1. 2001:838:1:1::1      0.0%  5  0.4  0.4  0.4  0.4
0.0 [.nl] Netherlands, The 12871 Concepts ICT
    ge-1-3-0.breda.ipv6.concepts-ict.net.
2. 2001:838:0:10::1    0.0%  5  2.3  2.4  2.3  2.4
0.0 [.nl] Netherlands, The 12871 Concepts ICT
3. 2001:7f8:1::a500:3257:1 0.0%  5  2.8  2.8  2.8  2.8  0.0
    ams-ix-1.ip.tiscali.net.
4. 2001:668:0:2::1:272  0.0%  5  2.9  3.0  2.9  3.0
0.0 [.eu] Europe Tiscali International Network B.V.
    so-6-1-0.ams22.ip6.tiscali.net.
5. 2001:668:0:2::1:481  0.0%  5  6.1  6.4  6.1  6.8
0.3 [.eu] Europe Tiscali International Network B.V.
    so-6-1-0.dus11.ip6.tiscali.net.
6. 2001:668:0:2::1:2    0.0%  5  9.5  9.4  9.2  9.6
0.2 [.eu] Europe Tiscali International Network B.V.
    so-0-1-0.fra40.ip6.tiscali.net.
7. 2001:668:0:2::1:292  0.0%  5  17.2 18.8 17.2 24.8
3.4 [.eu] Europe Tiscali International Network B.V.
    so-0-0-0.prg11.ip6.tiscali.net.
8. 2001:7f8:14::1:1     0.0%  5  18.0 20.1 18.0 25.4 3.1
    nix2-10ge.ipv6.cesnet.cz.
9. ???                  100.0  5  0.0  0.0  0.0  0.0  0.0

```

Can you post the output from traceroute6 from your Hardware Node / VE to 2001:838:1:1:210:dcff:fe20:7c7c (noc.sixxs.net) ?

Bye,
Thorsten
--
Thorsten Schifferdecker
tsd@debian.systs.org

Am Do, 24.01.2008, 21:47, schrieb Jan Tomasek:
> Hello Benoit,
>
>> Temporarily giving up messing with veth devices (forgot dhcpd server for
>> now), I concentrated on setting up a basic VS with both IPv4 and IPv6
>> addresses, using standard venet device.
>>
>> [...]
>
>> Result:
>>
>> - IPv4 VS works as expected (ping succeeds from hardware node, external
>> server in same subnet, external server on another subnet).
>>
>> - IPv6 VS doesn't work !!! (ping6 works *only* between VS and hardware
>> node, not from external sources; VS has IPv6 address correctly defined,

```
>> as per ifconfig).
>>
>> Am I missing something ?
>>
>> I suspected it may be useful to set sysctl
>> net.ipv6.conf.<iface>.forwarding to 1, but :
>> - this turns the box in "router" mode, disabling RA acceptance from
>> default router (annoying)
>> - this doesn't work either !!
>
> Very interesting! I'm having exactly same problem here. IPv6 was working
> perfectly for me after some problems before several months. Now I had to
> reboot my HW node and IPv6 is gone.
>
> When I try to ping VE from real world I'm getting:
>
>> semik@staff:~$ ping6 2001:718:1:e::23:3257
>> PING 2001:718:1:e::23:3257(2001:718:1:e::23:3257) 56 data bytes
>> From 2001:718:1:101::1 icmp_seq=0 Destination unreachable: Address
>> unreachable
>> From 2001:718:1:101::1 icmp_seq=3 Destination unreachable: Address
>> unreachable
>
> From HW node to VE or from one VE to another VE is IPv6 working fine. I
> think that HW node didn't sent some info to it's IPv6 gateway about
> presence of new IPv6 address. Sadly I'm clue less how to better debug
> this. But!
>
> I've discovered trick. Run alternate version of this on your HW node:
>
>> ifconfig eth0 add 2001:718:1:e::23:3257/64
>> sleep 5
>> ping6 -I 2001:718:1:e::23:3257 -c 1 2001:718:1:e::1
>> ifconfig eth0 del 2001:718:1:e::23:3257/64
>
> 2001:718:1:e::23:3257 - IPv6 address of VE
> 2001:718:1:e::1 - gateway
>
> That will cause that router will learn about presence of VE. It lasts
> for several hours. I discovered it at about 17 clock and now at 22 are
> VE still IPv6 online... but I'm afraid that clearing IPv6 neighbours
> table on router will kill this. Something is not working on HW node.
>
>
> I'm running 2.6.18-028stab051 on Debian Etch. Output of sysctl -a |grep
> net.ipv6 | sort >/tmp/chlivek.ipv6 is attached. Hope someone will point
> out where problem is. Yesterday I've reboot my second HW node...
>
```

> --
> -----
> Jan Tomasek aka Semik
> <http://www.tomasek.cz/>

Subject: Re: IPv6 ??
Posted by [Jan Tomasek](#) on Fri, 25 Jan 2008 07:12:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

Thorsten Schifferdecker wrote:
> @jan: how get your IPv6 connectivity ?

I've static setup on HW node:

```
iface eth0 inet6 static
    post-up /sbin/sysctl -w "net.ipv6.conf.eth0.autoconf=0"
    post-up /sbin/sysctl -w "net.ipv6.conf.all.forwarding=1"
    address 2001:718:1:e::23:3254
    gateway 2001:718:1:e::1
    netmask 64
```

Traceroute from HW node:

```
> chlivek:~# traceroute6 2001:838:1:1:210:dcff:fe20:7c7c
> traceroute to 2001:838:1:1:210:dcff:fe20:7c7c (2001:838:1:1:210:dcff:fe20:7c7c) from
2001:718:1:e::23:3254, 30 hops max, 16 byte packets
> 1 2001:718:1:e::1 (2001:718:1:e::1) 0.983 ms 4.869 ms 2.949 ms
> 2 2001:718:0:607::1 (2001:718:0:607::1) 4.943 ms 4.823 ms 2.895 ms
> 3 2001:718:0:607::2 (2001:718:0:607::2) 4.929 ms 4.867 ms 2.968 ms
> 4 ipv6-ge-2-1-0.prg11.ip.tiscali.net (2001:7f8:14::3:1) 3.359 ms 3.328 ms 4.439 ms
> 5 so-5-1-0.fra40.ip6.tiscali.net (2001:668:0:2::1:291) 9.113 ms 8.861 ms 14.915 ms
> 6 so-4-1-0.dus11.ip6.tiscali.net (2001:668:0:2::1:402) 12.322 ms 12.146 ms 15.893 ms
> 7 so-7-0-0.ams22.ip6.tiscali.net (2001:668:0:2::1:482) 15.88 ms 16.848 ms 15.919 ms
> 8 so-6-0-0.ams11.ip6.tiscali.net (2001:668:0:2::1:271) 15.889 ms 15.955 ms 14.971 ms
> 9 ge6-2-0.br0.ams3.nl.gbxs.net (2001:7f8:1::a500:9009:1) 15.511 ms 15.454 ms 15.972 ms
> 10 ams-ix2.ipv6.concepts.nl (2001:7f8:1::a501:2871:2) 15.437 ms 15.937 ms 15.837 ms
> 11 2001:838:0:14::2 (2001:838:0:14::2) 16.108 ms 15.294 ms 16.615 ms
> 12 2001:838:0:10::2 (2001:838:0:10::2) 21.508 ms 21.488 ms 17.428 ms
> 13 noc.sixxs.net (2001:838:1:1:210:dcff:fe20:7c7c) 21.508 ms 21.46 ms 18.441 ms
```

From VE (differnet IP, that 3257 was just experiemnt with my trick).

```
> root@idp:~# traceroute6 2001:838:1:1:210:dcff:fe20:7c7c
> traceroute to 2001:838:1:1:210:dcff:fe20:7c7c (2001:838:1:1:210:dcff:fe20:7c7c) from
```

```

2001:718:1:e::23:3248, 30 hops max, 16 byte packets
> 1 2001:718:1:e::23:3254 (2001:718:1:e::23:3254) 0.045 ms 0.015 ms 0.011 ms
> 2 2001:718:0:607::1 (2001:718:0:607::1) 1.105 ms 0.837 ms 1.381 ms
> 3 2001:718:0:607::2 (2001:718:0:607::2) 0.663 ms 0.44 ms 0.373 ms
> 4 ipv6-ge-2-1-0.prg11.ip.tiscali.net (2001:7f8:14::3:1) 0.566 ms 0.529 ms 0.529 ms
> 5 so-5-1-0.fra40.ip6.tiscali.net (2001:668:0:2::1:291) 8.551 ms 8.493 ms 8.499 ms
> 6 so-4-1-0.dus11.ip6.tiscali.net (2001:668:0:2::1:402) 11.302 ms 11.26 ms 11.256 ms
> 7 so-0-0-0.ams22.ip6.tiscali.net (2001:668:0:2::1:122) 15.086 ms 15.044 ms 15.039 ms
> 8 pos-7-5-0.ams11.ip6.tiscali.net (2001:668:0:2::630) 14.943 ms 14.945 ms 15.464 ms
> 9 ge6-2-0.br0.ams3.nl.gbxs.net (2001:7f8:1::a500:9009:1) 15.65 ms 15.52 ms 16.392 ms
> 10 ams-ix2.ipv6.concepts.nl (2001:7f8:1::a501:2871:2) 17.111 ms 15.812 ms 15.879 ms
> 11 2001:838:0:14::2 (2001:838:0:14::2) 16.031 ms 15.915 ms 15.91 ms
> 12 2001:838:0:10::2 (2001:838:0:10::2) 17.83 ms 18.28 ms 17.903 ms
> 13 noc.sixxs.net (2001:838:1:1:210:dcff:fe20:7c7c) 17.661 ms 18.263 ms 17.993 ms

```

Could you please try traceroute at 2001:718:1:e::23:3254? Path seems to be different from yours. But I'm not sure if that could indicate something bad in your IPv6 network setup. I'm not able to reach machines which are in our own network so it looks like something more fundamental. For IPv4 I will say that arp is not working but ...

--

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

Subject: Re: IPv6 ??
Posted by [Thorsten Schifferdeck](#) on Fri, 25 Jan 2008 07:45:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

@jan:

IPv6 traceroute from noc.sixxs.net @ SixXS NOC, AS12871 to 2001:718:1:e::23:3254 :

Hop	Node	Loss%	Sent	Last	Avg	Best	Worst
1.2	2001:838:1:1::1	0.0%	5	0.5	5.5	0.4	25.5
11.2	[.nl] Netherlands, The 12871 Concepts ICT ge-1-3-0.breda.ipv6.concepts-ict.net.						
2.0	2001:838:0:10::1	0.0%	5	2.6	3.0	2.4	4.5
0.9	[.nl] Netherlands, The 12871 Concepts ICT						
3.0	2001:7f8:1::a500:3257:1 ams-ix-1.ip.tiscali.net.	0.0%	5	3.1	2.9	2.7	3.2 0.2
4.0	2001:668:0:2::1:272	0.0%	5	3.2	3.2	3.1	3.3

```

0.1 [.eu] Europe      Tiscali International Network B.V.
    so-6-1-0.ams22.ip6.tiscali.net.
  5. 2001:668:0:2::1:121      0.0%  5  6.2  6.6  6.2  6.8
0.3 [.eu] Europe      Tiscali International Network B.V.
    so-4-0-0.dus11.ip6.tiscali.net.
  6. 2001:668:0:2::1:401      0.0%  5  9.8  9.6  9.3  9.8
0.2 [.eu] Europe      Tiscali International Network B.V.
    so-5-0-0.fra40.ip6.tiscali.net.
  7. 2001:668:0:2::1:292      0.0%  5 17.0 17.3 16.8 17.9
0.4 [.eu] Europe      Tiscali International Network B.V.
    so-0-0-0.prg11.ip6.tiscali.net.
  8. 2001:7f8:14::1:1          0.0%  5 18.6 41.7 18.2 134.9 52.1
    nix2-10ge.ipv6.cesnet.cz.
  9. 2001:718:1:e::23:3254     0.0%  5 17.9 17.8 17.4 18.0
0.2 [.cz] Czech Republic  CESNET Sub-T

```

work perfectly, btw, you can access the traceroute tool at <https://noc.sixxs.net/tools/traceroute/>, too

Bye,
 Thorsten
 --
 Thorsten Schifferdecker
 tsd@debian.systs.org

Am Fr, 25.01.2008, 08:12, schrieb Jan Tomasek:

```

> Hi,
>
> Thorsten Schifferdecker wrote:
>> @jan: how get your IPv6 connectivity ?
>
> I've static setup on HW node:
>
> iface eth0 inet6 static
>     post-up /sbin/sysctl -w "net.ipv6.conf.eth0.autoconf=0"
>     post-up /sbin/sysctl -w "net.ipv6.conf.all.forwarding=1"
>     address 2001:718:1:e::23:3254
>     gateway 2001:718:1:e::1
>     netmask 64
>
> Traceroute from HW node:
>
>> chlivek:~# traceroute6 2001:838:1:1:210:dcff:fe20:7c7c
>> traceroute to 2001:838:1:1:210:dcff:fe20:7c7c
>> (2001:838:1:1:210:dcff:fe20:7c7c) from 2001:718:1:e::23:3254, 30 hops
>> max, 16 byte packets
>>  1  2001:718:1:e::1 (2001:718:1:e::1)  0.983 ms  4.869 ms  2.949 ms

```

```

>> 2 2001:718:0:607::1 (2001:718:0:607::1) 4.943 ms 4.823 ms 2.895 ms
>> 3 2001:718:0:607::2 (2001:718:0:607::2) 4.929 ms 4.867 ms 2.968 ms
>> 4 ipv6-ge-2-1-0.prg11.ip.tiscali.net (2001:7f8:14::3:1) 3.359 ms
>> 3.328 ms 4.439 ms
>> 5 so-5-1-0.fra40.ip6.tiscali.net (2001:668:0:2::1:291) 9.113 ms
>> 8.861 ms 14.915 ms
>> 6 so-4-1-0.dus11.ip6.tiscali.net (2001:668:0:2::1:402) 12.322 ms
>> 12.146 ms 15.893 ms
>> 7 so-7-0-0.ams22.ip6.tiscali.net (2001:668:0:2::1:482) 15.88 ms
>> 16.848 ms 15.919 ms
>> 8 so-6-0-0.ams11.ip6.tiscali.net (2001:668:0:2::1:271) 15.889 ms
>> 15.955 ms 14.971 ms
>> 9 ge6-2-0.br0.ams3.nl.gbxs.net (2001:7f8:1::a500:9009:1) 15.511 ms
>> 15.454 ms 15.972 ms
>> 10 ams-ix2.ipv6.concepts.nl (2001:7f8:1::a501:2871:2) 15.437 ms
>> 15.937 ms 15.837 ms
>> 11 2001:838:0:14::2 (2001:838:0:14::2) 16.108 ms 15.294 ms 16.615 ms
>> 12 2001:838:0:10::2 (2001:838:0:10::2) 21.508 ms 21.488 ms 17.428 ms
>> 13 noc.sixxs.net (2001:838:1:1:210:dcff:fe20:7c7c) 21.508 ms 21.46 ms
>> 18.441 ms
>
>
> From VE (differnet IP, that 3257 was just experiemnt with my trick).
>
>> root@idp:~# traceroute6 2001:838:1:1:210:dcff:fe20:7c7c
>> traceroute to 2001:838:1:1:210:dcff:fe20:7c7c
>> (2001:838:1:1:210:dcff:fe20:7c7c) from 2001:718:1:e::23:3248, 30 hops
>> max, 16 byte packets
>> 1 2001:718:1:e::23:3254 (2001:718:1:e::23:3254) 0.045 ms 0.015 ms
>> 0.011 ms
>> 2 2001:718:0:607::1 (2001:718:0:607::1) 1.105 ms 0.837 ms 1.381 ms
>> 3 2001:718:0:607::2 (2001:718:0:607::2) 0.663 ms 0.44 ms 0.373 ms
>> 4 ipv6-ge-2-1-0.prg11.ip.tiscali.net (2001:7f8:14::3:1) 0.566 ms
>> 0.529 ms 0.529 ms
>> 5 so-5-1-0.fra40.ip6.tiscali.net (2001:668:0:2::1:291) 8.551 ms
>> 8.493 ms 8.499 ms
>> 6 so-4-1-0.dus11.ip6.tiscali.net (2001:668:0:2::1:402) 11.302 ms
>> 11.26 ms 11.256 ms
>> 7 so-0-0-0.ams22.ip6.tiscali.net (2001:668:0:2::1:122) 15.086 ms
>> 15.044 ms 15.039 ms
>> 8 pos-7-5-0.ams11.ip6.tiscali.net (2001:668:0:2::630) 14.943 ms
>> 14.945 ms 15.464 ms
>> 9 ge6-2-0.br0.ams3.nl.gbxs.net (2001:7f8:1::a500:9009:1) 15.65 ms
>> 15.52 ms 16.392 ms
>> 10 ams-ix2.ipv6.concepts.nl (2001:7f8:1::a501:2871:2) 17.111 ms
>> 15.812 ms 15.879 ms
>> 11 2001:838:0:14::2 (2001:838:0:14::2) 16.031 ms 15.915 ms 15.91 ms
>> 12 2001:838:0:10::2 (2001:838:0:10::2) 17.83 ms 18.28 ms 17.903 ms

```

>> 13 noc.sixxs.net (2001:838:1:1:210:dcff:fe20:7c7c) 17.661 ms 18.263
>> ms 17.993 ms
>
> Could you please try traceroute at 2001:718:1:e::23:3254? Path seems to
> be different from yours. But I'm not sure if that could indicate
> something bad in your IPv6 network setup. I'm not able to reach machines
> which are in our own network so it looks like something more
> fundamental. For IPv4 I will say that arp is not working but ...
>
>
>
> --
> -----
> Jan Tomasek aka Semik
> <http://www.tomasek.cz/>
>

Subject: Re: IPv6 ??
Posted by [Jan Tomasek](#) on Mon, 28 Jan 2008 17:37:11 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

problem was caused by setting:

```
# Controls which interfaces to send ARP requests and modify APR tables  
# on.  
NEIGHBOUR_DEVS=detect
```

this option arived with vzctl version 3.0.22 when I was upgrading from 3.0.18-1dso1.

After I removed that option and restarted HW node all IPv6 hosts get imediately online. Ufff! :)

--

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

Subject: Re: IPv6 ??
Posted by [kir](#) on Tue, 29 Jan 2008 09:09:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

Can you give us the the output of

/sbin/ip a l
and
fgrep ADDRESS /etc/vz/conf/*.conf
command run on your hardware node?

Perhaps the problem is that the network interface(s) that is/are present on the hardware node do not belong to the same IP subnets as your VEs.

Dmitry,

Maybe it makes sense to add a warning if vz_get_neighbour_devs() returns an empty list, with the URL to a wiki page with long explanation?

Jan Tomasek wrote:

> Hi,
>
> problem was caused by setting:
>
> # Controls which interfaces to send ARP requests and modify APR tables
> # on.
> NEIGHBOUR_DEVS=detect
>
> this option arived with vzctl version 3.0.22 when I was upgrading from
> 3.0.18-1dso1.
>
> After I removed that option and restarted HW node all IPv6 hosts get
> imediately online. Ufff! :)
>

Subject: Re: IPv6 ?? *SOLVED*

Posted by [Benoit Branciard](#) on Wed, 30 Jan 2008 15:32:15 GMT

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

>
> - IPv6 VS doesn't work !!! (ping6 works *only* between VS and hardware
> node, not from external sources; VS has IPv6 address correctly defined,
> as per ifconfig).
>

For list archives:

It was caused by a small bug in vzctl scripts which handle "NEIGHBOUR_DEVS=detect" feature introduced in vzctl 3.0.22 : neighbour devices were not detected for IPv6.

Fix is here:

<http://git.altlinux.org/people/ldv/packages/?p=vzctl.git;a=commitdiff;h=maint>

Thanks to Jan, Kir and Dmitry !

> I suspected it may be useful to set sysctl
> net.ipv6.conf.<iface>.forwarding to 1, but :
> - this turns the box in "router" mode, disabling RA acceptance from
> default router (annoying)

IPv6 forwarding is still mandatory, but this implies manual configuration of IP address, netmask and default route on hardware node. This last point may be annoying: in our case the default route is hardware-dependent, so if some day we need to replace the router (or the router interface), default route will become invalid and manual reconfiguration of all openVZ hardware nodes will be needed. Too bad Linux kernel doesn't allow enabling forwarding without disabling autoconfiguration.

--

Ce message a ete verifie par MailScanner pour des virus ou des polluriels et rien de suspect n'a ete trouve.

Subject: Re: IPv6 ?? *SOLVED*

Posted by [Jan Tomasek](#) on Thu, 31 Jan 2008 07:45:16 GMT

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

Benoit Branciard wrote:

>> I suspected it may be useful to set sysctl
>> net.ipv6.conf.<iface>.forwarding to 1, but :
>> - this turns the box in "router" mode, disabling RA acceptance from
>> default router (annoying)
>
> IPv6 forwarding is still mandatory, but this implies manual
> configuration of IP address, netmask and default route on hardware node.
> This last point may be annoying: in our case the default route is
> hardware-dependent, so if some day we need to replace the router (or the
> router interface), default route will become invalid and manual
> reconfiguration of all openVZ hardware nodes will be needed.
> Too bad Linux kernel doesn't allow enabling forwarding without disabling
> autoconfiguration.

I'm not sure, but I think that my setup is router independent. For my virtual servers I got prefix 2001:718:1:e::/64 (which is unbelievable waste of IP range but they say this is normal in IPv6). This network block is assigned to dedicated VLAN. IPv6 Gateway for all physical

systems in that network is 2001:718:1:e::1.

My setup on HW node is:

```
iface eth0 inet6 static
    post-up /sbin/sysctl -w "net.ipv6.conf.eth0.autoconf=0"
    post-up /sbin/sysctl -w "net.ipv6.conf.all.forwarding=1"
    address 2001:718:1:e::23:3254
    gateway 2001:718:1:e::1
    netmask 64
```

I'm trying to stay away of MAC dependent automatically assigned IPs because of dependence on hardware.

I hope this might you, if needed I might try to ask our network specialists about more details...

Best regards

--

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

Subject: Re: IPv6 ?? *SOLVED*
Posted by [Benny Amorsen](#) on Thu, 31 Jan 2008 09:01:27 GMT
[View Forum Message](#) <> [Reply to Message](#)

Jan Tomasek <jan@tomasek.cz> writes:

> I'm not sure, but I think that my setup is router independent. For my
> virtual servers I got prefix 2001:718:1:e::/64 (which is unbelievable
> waste of IP range but they say this is normal in IPv6).

Don't worry too much about address waste. E.g. every single IPv4 address has an equivalent /48 with 6to4.

Also, the current allocation policies only apply to 1/8 of all IPv6 space. It is not expected that the other 7/8 will ever be needed -- they are reserved in case the architects were wrong and we need to be more frugal.

/Benny

Subject: Re: IPv6 ?? *SOLVED*

> Benoit Branciard wrote:

>>

>> IPv6 forwarding is still mandatory, but this implies manual
>> configuration of IP address, netmask and default route on hardware
>> node. This last point may be annoying: in our case the default route
>> is hardware-dependent, so if some day we need to replace the router
>> (or the router interface), default route will become invalid and
>> manual reconfiguration of all openVZ hardware nodes will be needed.
>> Too bad Linux kernel doesn't allow enabling forwarding without
>> disabling autoconfiguration.

>

> I'm not sure, but I think that my setup is router independent. For my
> virtual servers I got prefix 2001:718:1:e::/64 (which is unbelievable
> waste of IP range but they say this is normal in IPv6). This network
> block is assigned to dedicated VLAN. IPv6 Gateway for all physical
> systems in that network is 2001:718:1:e::1.

Here's the clue: you have been assigned a static gateway IP address by
your network provider. So no problem.

In our case we are operating our network, and have chosen to not set
static gateway IP addresses, relying on autoconfiguration to discover the
EUI-64 (MAC-dependant) default route IP. That sounded nice (nothing to
hardcode) until openVZ came on.

Maybe we could change our router config policy, and assign static IP
addresses to router interfaces. But that's not trivial operation,
involving network downtime, and we won't do that without motivated reasons.

--

Ce message a ete verifie par MailScanner
pour des virus ou des polluriels et rien de
suspect n'a ete trouve.
