
Subject: Assign subnet to VM? Or increase routing table limits?

Posted by [tywe](#) on Sat, 28 Aug 2010 03:32:00 GMT

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

Hello,

I'm trying to find a way to add a subnet of IPs to a VM instead of having to add individual IP addresses one by one. This is for IPv6, so it's quite common for clients to request larger subnets, but I always force them to have a block of 64 individual addresses or similar amount (128, 256, etc).

However, I recently hit a wall in the routing table around 1200 or so entries where I start getting this error when trying to add more IPv6 addresses:

RTNETLINK answers: Cannot allocate memory

And yes, the server has plenty of free RAM available, over 10GB is free.

So, I tried to find info on how to increase the routing table limits, but didn't have much luck there, so then I started trying to figure out a way to add a subnet instead of individual IP addresses, since this would be much easier for the routing table to handle. OpenVZ only seemed to allow individual IPs to be added, so I tried setting up a route for the subnet to the VPS without any luck. The route existed and the VM had everything setup on that end with a few IPs binding to it, so it seemed like it should work, but doesn't, which I guess is due to something internal with OpenVZ where a typical route won't be enough to complete the connection.

I also tried messing with veth devices instead of venet, but couldn't get anything to route through them either, which may have something to do with the original config and me not being able to take this server down to start testing any major network configurations.

Anyhow, hopefully someone can point me to some sort of solution for any of these issues and be able to add a subnet of IP addresses? If adding a subnet isn't possible for some reason, then hopefully we can at least increase my routing table limits?

Thanks very much for any suggestions you may have!

-tywe

Subject: Re: Assign subnet to VM? Or increase routing table limits?

Posted by [tywe](#) on Mon, 30 Aug 2010 19:56:25 GMT

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

Anyone have any ideas at all on this? Would maybe upgrading my kernel improve anything? I'm currently running 2.6.18-164.15.1.el5.028stab068.9 and need to upgrade anyhow, but have been trying to prolong it if possible to prevent any new problems or complaints from clients.

Thanks again.

Subject: Re: Assign subnet to VM? Or increase routing table limits?

Posted by [tywe](#) on Thu, 09 Sep 2010 22:54:21 GMT

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

Sorry to reply to myself again, but I'm really hoping someone has an idea to resolve this? It's hard to believe no one else has ever ran into any routing table limits at around 1200 entries when doing VPS hosting, especially since we don't seem to be able to assign subnets and have to do each IP individually.

If no one else has ever ran into this, then I expect it'll be happening soon, since IPv6 is getting more popular and IPv4 will run out soon.

We urgently need one of the following solutions:

- 1) Allow subnets to be assigned instead of forcing individual IPs (or tell me what I'm doing wrong if this is already allowed)
- 2) Increase routing table limit on host node to allow much more than 1200 entries. (or tell me why mine is running out if this shouldn't be an issue)

Thanks

Subject: Re: Assign subnet to VM? Or increase routing table limits?

Posted by [lars.bailey](#) on Fri, 07 Jan 2011 04:09:33 GMT

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

I'm actually surprised that you haven't gotten some kind of response to this thread.
Using this example;

Prefix/L: fd
Global ID: 2a92636461
Subnet ID: 0491
Combined/CID: fd2a:9263:6461:0491::/64
IPv6 addresses: fd2a:9263:6461:0491:xxxx:xxxx:xxxx:xxxx

The subnet of this prefix is;

0491

In other words, this is the end-point.

Any sub-netting schema for address management, should have been done somewhere in your router topology.

This is a matrix of a IPv6 prefix.

FD2A:9263:6461:0491:0000:0000:0000:0000
XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX

```

||| ||| ||| ||| ||| ||| |||
||| ||| ||| ||| ||| ||| |||128
||| ||| ||| ||| ||| ||| |||124
||| ||| ||| ||| ||| ||| |||120
||| ||| ||| ||| ||| ||| |||116
||| ||| ||| ||| ||| ||| |||112
||| ||| ||| ||| ||| ||| |||108
||| ||| ||| ||| ||| ||| |||104
||| ||| ||| ||| ||| ||| |||100
||| ||| ||| ||| ||| ||| |||96
||| ||| ||| ||| ||| ||| |||92
||| ||| ||| ||| ||| ||| |||88
||| ||| ||| ||| ||| ||| |||84
||| ||| ||| ||| ||| ||| |||80
||| ||| ||| ||| ||| ||| |||76
||| ||| ||| ||| ||| ||| |||72
||| ||| ||| ||| ||| ||| |||68
||| ||| ||| ||| ||| ||| |||64
||| ||| ||| ||| ||| ||| |||60
||| ||| ||| ||| ||| ||| |||56
||| ||| ||| ||| ||| ||| |||52
||| ||| ||| ||| ||| ||| |||48
||| ||| ||| ||| ||| ||| |||44
||| ||| ||| ||| ||| ||| |||40
||| ||| ||| ||| ||| ||| |||36
||| ||| ||| ||| ||| ||| |||32
||| ||| ||| ||| ||| ||| |||28
||| ||| ||| ||| ||| ||| |||24

```

What you would have to do, it cut the /64, into a /112.

```

FD2A:9263:6461:0491:0000:0000:0000::/112
FD2A:9263:6461:0491:0000:0000:0001::/112
FD2A:9263:6461:0491:0000:0000:0002::/112
FD2A:9263:6461:0491:0000:0000:0003::/112
FD2A:9263:6461:0491:0000:0000:0004::/112
FD2A:9263:6461:0491:0000:0000:0005::/112
FD2A:9263:6461:0491:0000:0000:0006::/112
=>

```

Assign the Node server, an IP from;

```
FD2A:9263:6461:0491:0000:0000:0000::1A
```

If you are using non-bridged VETH, each requires an IP from a different subnet.

A VETH from container VE101 would get;

FD2A:9263:6461:0491:0000:0000:0001:1A

The VE would get;

FD2A:9263:6461:0491:0000:0000:0001:65 "65 is 101 in binary"

All that is needed, is a route to the Node server's source-route interface, via the VETH IP, in a container.

But there is a caveat, to this type of setup.

Never use the IP command, to list IPv6 neighbors.

For whatever reason, it blows away the routing, to where you are back to square one. (this has happened to me)

For what it's worth, I think Ethernet bridging is really the solution for this type of setup. (one aggregate prefix via bridge)

Hope it helps.

Subject: Re: Assign subnet to VM? Or increase routing table limits?

Posted by [tywe](#) on Fri, 07 Jan 2011 06:27:30 GMT

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

Hi and thanks very much for the reply! I was beginning to think no one was going to help around here on either thread I've started asking for help, lol

I'm a bit confused on exactly how to set it all up the way you described, but your reply was very informative and definitely has pointed me in the right direction, so as soon as I have some time to spend on this, I'll do some testing and will report back on how it goes.

Thanks again!

Subject: Re: Assign subnet to VM? Or increase routing table limits?

Posted by [lars.bailey](#) on Thu, 13 Jan 2011 08:09:17 GMT

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

I posted a thread on using IPv6 with OVZ, in two arenas.

One on auto-configuration and its issues, and one on static IP, and some things that can go wrong.

Most involve improper subnetting schemas.

This post;

http://forum.openvz.org/index.php?t=tree&goto=41388&#msg_41388

which I hope,will head you in the right direction.(read the first message)
Don't take to heart the first one,as it actually represents a bad example,and see if you may be trying something along that line.(adding un-necessary routes due to invalid subnets)
The second message,explains in detail,a proper subnetting example,and how I configured IPv6 for testing.
The long and short for us now,was to fit each Node with Quagga.
This was mainly for persistent IPv6 configurations,even though,we do use OSPF.(the Nodes are now ABR routers)
If the post isn't what you are doing/trying,keep the thread going,and I'll do what I can,to help.

Regards

Subject: Re: Assign subnet to VM? Or increase routing table limits?

Posted by [lars.bailey](#) on Thu, 13 Jan 2011 08:32:28 GMT

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

P.S

Replies here,are a rarity.
Its a shame actually.
Do not use VENET with IPv6,as there is a doc on the WIKI,that explains the issues with IPv6 and VENET.
My example from the link I gave you,uses virtual Ethernet.(VETH)
If you have never used virtual Ethernet,the VETH interface,is basically a VE's gateway.
Which means,a VE only requires a default route.
Using a firewall program like Shorewall,you can "zone" the VETH,
and condition what parameters you allow the container,including traffic shaping.
I would also recommend that you download the subnetting tool I use.

ipv6gen.pl

Not only will it give you a proper subnetting schema,you will find out,just how many host IPv6 addresses exist,in a given prefix length.
To give an example of a /116;

FD22:A075:AFD0:E096:0000:0001:0000:1000/116

you would have IPv6 address range from;

FD22:A075:AFD0:E096:0000:0001:0000:1001/116

=>

FD22:A075:AFD0:E096:0000:0001:0000:1FFF/116

I think you will find you answer, in your subnetting schema.
Keep me posted.
