

---

Subject: \*SOLVED\* VLAN on VETH device?  
Posted by [szocske](#) on Mon, 25 Sep 2006 09:31:20 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hi!

I'd like to issue vconfig commands within a VE to set up VLANs on its veth interface. (which it beleives to be eth0.) But there is no /proc/net/vlan directory inside the VEs.  
I have the 8021q module loaded in the kernel, so vconfig works fine outside, in the HW node, on the real eth0.

Is this a known limitation, a bug, or my misconfiguration?

OpenVz patch version 026test015.

---

---

Subject: Re: VLAN on VETH device?  
Posted by [szocske](#) on Tue, 26 Sep 2006 09:00:02 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

I'd hate to rush things guys, but this is a serious roadblock in our openVZ deplolyment. We might even consider purchasing a bunch of closedVZ licenses, if that solves this particular issue. (hard to imagine how it could be done outside the kernel and its GPL though...)

Assuming there is no support yet:

How big of an effort would you estimate the development of VLAN support on veth devices? Is there some theoretical problem with it, or just omething noone ever bothered to do just yet?

Is it something we could undertake (and contribute back of course) or is it possible to contract you guys for such targeted development? (we are an open source shop, would prefer the feature in the OS line.)

---

---

Subject: Re: VLAN on VETH device?  
Posted by [Andrey Mirkin](#) on Tue, 26 Sep 2006 11:29:15 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

vlan proc entry is created only in VE0 as vlan is not virtualized, so you can't see it (/proc/net/vlan) in VE.

You can create vlan interface in VE0 and move it to VPS with help of  
vzctl set --netdev\_add <name>  
command.

---

---

Subject: Re: VLAN on VETH device?  
Posted by [szocske](#) on Tue, 26 Sep 2006 11:47:21 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Great idea, thanks.

Not the full solution we are looking for though

Is it possible to --netdev\_add the same tagging interface to multiple VEs at the same time? Can they then communicate through it? Also, the packets sent out on such an interface would contain the MAC address of the real eth0, not the one assigned to the veth, right?

Anyway, we are interested in estimating the effort needed to make VLANs work from VEs: our software controls VLANs directly with ioctl, and we'd like to be able to run it in VEs.

---

---

Subject: Re: VLAN on VETH device?

Posted by [Andrey Mirkin](#) on Tue, 26 Sep 2006 13:39:07 GMT

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

---

You can create VLAN in VE0 and add it to bridge with other veth devices, so VPSes which are added to this bridge will be able to communicate with each other.

See [http://wiki.openvz.org/Virtual\\_Ethernet\\_device#Virtual\\_ethernet\\_devices\\_.2B\\_VLAN](http://wiki.openvz.org/Virtual_Ethernet_device#Virtual_ethernet_devices_.2B_VLAN) for details.

You are right, the packets sent out on such an interface would contain the MAC address of the real eth0.

We need to analyze vlan module to estimate the time we will need to virtualize it.

---

---

Subject: Re: VLAN on VETH device?

Posted by [szocske](#) on Thu, 02 Nov 2006 08:59:58 GMT

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

---

Andrey Mirkin wrote on Tue, 26 September 2006 09:39

We need to analyze vlan module to estimate the time we will need to virtualize it.

Hi there!

Any news on this?

We are interested in sponsoring or performing the development of vlan support on veth devices.

We migrated most of our test servers from Xen and VMware to OpenVZ in the recent months. Since we tuned the resource quotas high enough, our developers are happy with it. (well, they are happy to have more servers, and happy not to have to care what goes on below. except for vlan-related tests )

Thanks for the great tool anyhow!

---

---

Subject: Re: VLAN on VETH device?

Posted by [Andrey Mirkin](#) on Fri, 03 Nov 2006 16:22:50 GMT

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

---

Hello!

Vlan support inside VE will be available soon. I have almost finished patch. After some testing I'll put it here.

---

---

Subject: Re: VLAN on VETH device?

Posted by [szocske](#) on Fri, 03 Nov 2006 22:22:35 GMT

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

---

Andrey Mirkin wrote on Fri, 03 November 2006 11:22

Vlan support inside VE will be available soon. I have almost finished patch. After some testing I'll put it here.

This is fantastic! We are looking forward to it. I am truly grateful for the effort.

---

---

Subject: Re: VLAN on VETH device?

Posted by [Andrey Mirkin](#) on Tue, 07 Nov 2006 16:19:50 GMT

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

---

Hello!

The patch for 2.6.18 kernel is ready.

You can try it right now.

vlan virtualization will be available in next kernel release.

#### File Attachments

1) [diff-net-vlan-virt](#), downloaded 489 times

---

---

Subject: Re: VLAN on VETH device?

Posted by [szocske](#) on Wed, 08 Nov 2006 10:49:47 GMT

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

---

Andrey Mirkin wrote on Tue, 07 November 2006 11:19Hello!

The patch for 2.6.18 kernel is ready.

You can try it right now.

You guys are simply great. We'll give it a try as soon as we can. (resource allocation issues, both human and computational, prohibit me from trying right away, I'll probably end up compiling it over

the weekend.)

---

Subject: Re: VLAN on VETH device?

Posted by [szocske](#) on Thu, 09 Nov 2006 10:19:57 GMT

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

---

Andrey Mirkin wrote on Tue, 07 November 2006 11:19: The patch for 2.6.18 kernel is ready.  
You can try it right now.  
vlan virtualization will be available in next kernel release.

We have it working, and see no regression in regular functionality.  
/proc/net/vlan directory exists and can be populated via vconfig.

But any vlan subinterface created in one VE is visible in all (the veth device is called eth0 in all VEs)

ifconfig -a does not list "foreign" vlans.

It is not possible to create vlans with the same vlan ID in multiple VEs:

```
eserver25-vz6:~# vconfig add eth0 106
Added VLAN with VID == 106 to IF -:eth0:-
```

```
eserver25-vz7:~# vconfig add eth0 107
Added VLAN with VID == 107 to IF -:eth0:-
```

```
eserver25-vz6:~# vconfig add eth0 107
ERROR: trying to add VLAN #107 to IF -:eth0:- error: Invalid argument
```

```
eserver25-vz6:~# cat /proc/net/vlan/
config  eth0.106 eth0.107
eserver25-vz6:~# cat /proc/net/vlan/config
VLAN Dev name | VLAN ID
Name-Type: VLAN_NAME_TYPE_RAW_PLUS_VID_NO_PAD
eth0.106      | 106 | eth0
eserver25-vz6:~# ls /proc/net/vlan/
config eth0.106 eth0.107
eserver25-vz6:~# cat /proc/net/vlan/eth0.107
eth0.107 VID: 107 REORDER_HDR: 1 dev->priv_flags: 1
    total frames received      0
    total bytes received      0
Broadcast/Multicast Rcvd      0

    total frames transmitted    0
    total bytes transmitted    0
    total headroom inc          0
```

```
total encap on xmit      0
Device: eth0
INGRESS priority mappings: 0:0  1:0  2:0  3:0  4:0  5:0  6:0  7:0
EGRESSSS priority Mappings:
```

Can we have this revised please?

Editing:

Just noticed, same goes for host machine:

```
eserver25:~# ls /proc/net/vlan/
config eth0.106 eth0.107
eserver25:~# cat /proc/net/vlan/config
VLAN Dev name   | VLAN ID
Name-Type: VLAN_NAME_TYPE_RAW_PLUS_VID_NO_PAD
```

Editing again:

And I can't shut down the VE which has a VLAN configured:

```
eserver25:~# vzctl stop 109
Stopping VPS ...
```

```
Message from syslogd@localhost at Thu Nov  9 11:34:37 2006 ...
localhost kernel: unregister_netdevice: waiting for eth0 to become free. Usage count = -1
```

```
Message from syslogd@localhost at Thu Nov  9 11:35:17 2006 ...
localhost last message repeated 4 times
```

Meanwhile, the vlans disappear from the remaining machines:

```
eserver25-vz6:~# ls /proc/net/vlan/
config
eserver25-vz6:~# cat /proc/net/vlan/config
VLAN Dev name   | VLAN ID
Name-Type: VLAN_NAME_TYPE_RAW_PLUS_VID_NO_PAD
eth0.106       | 106 | eth0
```

---

Subject: Re: VLAN on VETH device?

Posted by [Andrey Mirkin](#) on Thu, 09 Nov 2006 16:29:00 GMT

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

---

Hello!

The previous patch was not full, sorry.

Please try this patch. This patch still can be not very stable, but with you help and feedback we will try to fix all issues.

## File Attachments

1) [diff-net-vlan-virt-20061109](#), downloaded 459 times

---

---

Subject: Re: \*SOLVED\* VLAN on VETH device?

Posted by [szocske](#) on Thu, 16 Nov 2006 09:15:49 GMT

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

---

Hi again!

Sorry for the high latency, we had to provision the new sun servers we purchased for openvz

The updated patch works, we can now ping over a vlan from one VE to another. (tested within a HW node only)

Stopping VEs works fine too.

/proc/net/vlan looks funny though:

```
sun03:/# ls /proc/net/vlan
```

```
config eth0.42 eth0.42 eth2.42
```

VLAN ID 42 was configured on the veth-s of two VEs (both eth0) and the HW node eth2. Looks the same from HW node and VEs.

We'll continue testing across a switch, and with our ioctls, and report.

One more (probably unrelated) quirk:

Despite setting up /etc/hostname correctly in the VEs, when started, `hostname` returns the HW node hostname in all of them.

Better yet: issuing the hostname command within a VE to change the hostname changes it for all VEs, and the HW node as well!

Used to work fine on 2.6.16.

Debian stable + vlan-enabled kernel

---

---

Subject: Re: \*SOLVED\* VLAN on VETH device?

Posted by [arpad](#) on Fri, 17 Nov 2006 09:21:00 GMT

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

---

Hi!

I'm working with szocske.

There is another problem with the OpenVZ maybe it's related to the hostname problem. When I stop a running VE I got this in the dmesg:

```
BUG: warning at fs/proc/generic.c:916/__remove_proc_entry()
[<c01b832a>] __remove_proc_entry+0xea/0x100
[<c01b837b>] remove_proc_loc_entry+0x3b/0x50
[<c01b83ff>] remove_proc_entry+0x1f/0x60
[<f8f8c44d>] fini_ve_sysctl+0x2d/0x40 [vzmon]
[<f8f90ba2>] env_cleanup+0xf2/0x1a0 [vzmon]
[<f8f90ce6>] vzmond_helper+0x46/0x60 [vzmon]
[<f8f90ca0>] vzmond_helper+0x0/0x60 [vzmon]
[<c01012f5>] kernel_thread_helper+0x5/0x10
```

We are using Debian patched kernels (it's our company policy) and the basic Debian patch-remade-2.6.18-openvz-debian patch for OpenVZ. We applied yours new VLAN patch too.

I hope it helps. If you need we can compile a pure Vanilla kernel for debug purposes.

Thanks for any help!

---

Subject: Re: **\*SOLVED\*** VLAN on VETH device?  
Posted by [arpad](#) on Fri, 24 Nov 2006 08:28:30 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hi!

I'm here again. We compiled a vanilla 2.6.18 kernel with patch-ovz028test005.1-combined and the patch you generously provided for VLAN. The only problem was that one file had moved to another directory so I had to rewrite its location in the patch.

The 802.1Q tagging seems to work but after one day usage we noticed that our functional testing not running anymore because the ssh daemon running in the VE randomly dropped the connections. I looked into the logs but found nothing relevant. I checked the quotas but none of them was violated.

So if you have any ideas I will be very happy. We need to use 802.1Q for our tests and we must keep running another machine with Xen just for this.

Thanks,

Arpad

---

Subject: Re: **\*SOLVED\*** VLAN on VETH device?  
Posted by [Andrey Mirkin](#) on Fri, 24 Nov 2006 12:54:43 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Vlan virtualization will be available in 2.6.18-028test006 kernel.

---

---

Subject: Re: \*SOLVED\* VLAN on VETH device?  
Posted by [szocske](#) on Fri, 24 Nov 2006 13:34:01 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hi!

Arpad and I are working together on putting your diff-net-vlan-virt patch to testing.  
Are you saying we should wait for 2.6.18-028test006 and report issues only with that version?

---

---

Subject: Re: \*SOLVED\* VLAN on VETH device?  
Posted by [Andrey Mirkin](#) on Fri, 24 Nov 2006 14:10:09 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

You can wait for 2.6.18-028test006 kernel or can use new patch from git:  
<http://git.openvz.org/?p=linux-2.6.18-openvz;a=commit;h=f629dec5ea2d32d05ccedcb96c47e41971b6200c>

---