

---

Subject: [ SOLVED ] Pb with multiple identical nfs mount dir in different VEs  
Posted by [mistur](#) on Wed, 16 Apr 2008 08:50:44 GMT

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

---

Hi,

I have a problem with an identical nfs mount in different VEs

for example, I have inside some VEs :

```
$ df /mnt/data
10.0.0.1:/nfs/data 134148096 43458560 90689536 33% /mnt/data
```

if I perform a reboot on one VE, on all others VEs where this partition is mounted, I got :

```
$ df /mnt/data
10.0.0.1:/nfs/data - - - - /mnt/data
```

so the partition is still mount but impossible to do an access on.

at the same time I got this in kernel log on the host :

```
[258012.479557] unregister_vlan_device: ERROR: Tried to remove a non-vlan device with VLAN
code, name: eth1.35 priv_flags: 0
[258012.646468] unregister_vlan_device: ERROR: Tried to remove a non-vlan device with VLAN
code, name: eth1.32 priv_flags: 0
[258013.466439] nfs: RPC call returned error 101
[258013.711497] nfs: RPC call returned error 101
[258013.715907] nfs: RPC call returned error 101
```

As soon as I removed all references of that mount point in all VEs. I can remount it, it's working.

VE 201 :

```
NETIF=" ifname=eth1.32,mac=00:16:3E:02:01:1E,host_ifname=veth201.32,
host_mac=FE:FF:FF:FF:FF:FF;ifname=eth1.35,mac=00:16:3E:02:01
:1E,host_ifname=veth201.35,host_mac=FE:FF:FF:FF:FF:FF;ifname
=eth1.36,mac=00:16:3E:02:01:1E,host_ifname=veth201.36,host_m ac=FE:FF:FF:FF:FF:FF
"NETIF=" ifname=eth1.32,mac=00:16:3E:02:02:1E,host_ifname=veth202.32,
host_mac=FE:FF:FF:FF:FF:FF;ifname=eth1.35,mac=00:16:3E:02:02
:1E,host_ifname=veth202.35,host_mac=FE:FF:FF:FF:FF:FF;ifname
=eth1.36,mac=00:16:3E:02:02:1E,host_ifname=veth202.36,host_m ac=FE:FF:FF:FF:FF:FF "
```

VE 202 :

```
NETIF=" ifname=eth1.32,mac=00:16:3E:02:02:1E,host_ifname=veth202.32,
host_mac=FE:FF:FF:FF:FF:FF;ifname=eth1.35,mac=00:16:3E:02:02
:1E,host_ifname=veth202.35,host_mac=FE:FF:FF:FF:FF:FF;ifname
=eth1.36,mac=00:16:3E:02:02:1E,host_ifname=veth202.36,host_m ac=FE:FF:FF:FF:FF:FF "
```

Ubuntu Dapper 6.06 LTS up to date

```
$ uname -a
```

```
Linux vz01 2.6.18-028stab053 #1 SMP Wed Feb 20 12:15:30 CET 2008 x86_64 GNU/Linux
```

```
$ vzctl --version
```

```
vzctl version 3.0.22
```

if someone has a idea...

thanks

Yoann Moulin

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs

Posted by [mistur](#) on Thu, 17 Apr 2008 13:18:55 GMT

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

no idea ?

If you need more information, I can compile new kernel with debug option but I don't know which one can be usefull

or I'm able to do lot's of test

attach : config-2.6.18-028stab053

regards

Yoann

#### File Attachments

1) [config-2.6.18-028stab053](#), downloaded 470 times

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs

Posted by [maratrus](#) on Mon, 21 Apr 2008 13:48:23 GMT

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

Hello,

Sorry for delay.

1.  
Could you possibly specify how did you exactly configure network inside VE and on HN? You can

show us the exact commands if it is convenient for you.

2.

How do you mount content via NFS?  
Something like that? <http://wiki.openvz.org/NFS>

Thank You!

---

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs  
Posted by [mistur](#) on Tue, 22 Apr 2008 09:45:05 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

maratrus wrote on Mon, 21 April 2008 15:48Hello,

Sorry for delay.  
that's ok

maratrus wrote on Mon, 21 April 2008 15:481.

Could you possibly specify how did you exactly configure network inside VE and on HN? You can show us the exact commands if it is convenient for you.

2.

How do you mount content via NFS?  
Something like that? <http://wiki.openvz.org/NFS>

see in attachement for all details about networks and nfs configuration

I can do all tests needed

Thanks for your help

Yoann

### File Attachments

---

- 1) [host.txt](#), downloaded 618 times
  - 2) [ve.txt](#), downloaded 547 times
- 

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs  
Posted by [maratrus](#) on Wed, 23 Apr 2008 11:46:26 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hello,

I've been trying to reproduce this issue but I haven't been able to do that.

This sounds like a bug.

Is that issue reproducible on any OS template? Please try to do the same with VE based on the non ubuntu template for example with centos or fedora template.

Is that HN is a production node? Could we get an access to it if we won't be able to reproduce this behavior?

P.S.

Quote:As soon as I removed all references of that mount point in all VEs. I can remount it, it's working.

Could you please clarify what does it mean?

Quote:if I perform a reboot on one VE,

vzctl restart \$VE\_ID?

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs

Posted by [mistur](#) on Tue, 29 Apr 2008 14:27:14 GMT

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

---

after a couple of day of investigation, I have more information about the problem.

In my test, I have 3 VE, one centos5 i386, one centos5 x86\_64, one ubuntu Dapper 6.06

in all VE I have an ip set on eth1.35 and all have the same nfs directory mounted.

If I run a ifdown eth1.35 on centos, no problem but I if run ifdown eth1.35 on the ubuntu, nfs mount point borked... until I run an ifup eth1.35.

So it seem that the script inside the VE on ubuntu is doing something that not supposed to do and change something inside vlan configuration for others VE.

in attachment the result of

```
strace ifdown eth1.35 > /tmp/ifdown.eth1.35.trace 2>&1
```

Yoann

---

### File Attachments

1) [ifdown.eth1.35.trace](#), downloaded 554 times

---

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs

Posted by [mistur](#) on Tue, 29 Apr 2008 15:46:41 GMT

Some additional informations usefull I guess :

for each VE, I have 3 mount points over the same vlan 35; but each are mounted with different nfs server IP .

IP1:/mnt/data1  
IP2:/mnt/data2  
IP3:/mnt/data3

IP1 IP2 IP3 are on the same vlan and all of them are set on the same nfs server for now (can be move separately, each IP is associate to a drbd device)

here come the strange part of this bug :

only one mount point is affected by this bug, the 2 others are working perfectly in others VEs after set down eth1.35 inside the VE running ubuntu.

edit: the vlan35 is also the network used for root over nfs for the HN, and where are hosted /vz/private for VE.

rootfs for HN and VE are on a separate nfs server dedicated to rootfs and not on the same server used for /mnt/data1, /mnt/data2 and /mnt/data3

the HN is completely diskless

but, it's still complicated to fully qualified this bug...

Yoann

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs  
Posted by [maratrus](#) on Wed, 30 Apr 2008 08:07:22 GMT

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

---

Hi,

I've suspected that something is not quite clear with ubuntu "ifdown\*" utilities.  
But the fact is that in redhat based OS ifdown is a simple script so it's easy to investigate it:  
it calls appropriate script in /etc/sysconfig/network-scripts/.  
And the only thing I've found there is:

```
if [ -n "$VLAN" -a -x /sbin/vconfig ]; then
    # 802.1q VLAN
    if echo ${DEVICE} | LANG=C egrep -v '(:' | LANG=C egrep -q
'(eth|bond)[0-9][0-9]*\.[0-9][0-9]?[0-9]?[0-9]?' \
    || echo ${DEVICE} | LANG=C egrep -q 'vlan[0-9][0-9]?[0-9]?[0-9]?' ; then
        [ -f /proc/net/vlan/${DEVICE} ] && {
```

```
    /sbin/vconfig rem ${DEVICE}
}
fi
fi
```

So,

1. If it's possible we can investigate src files of appropriate utilities in ubuntu OS.

2.

Quote:

```
strace ifdown eth1.35 > /tmp/ifdown.eth1.35.trace 2>&1
```

ifdown can generate another process so could you possible try to it in the following way:

```
strace -fF -o out.txt ifdown eth1.35
```

Thank You!

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs

Posted by [mistur](#) on Wed, 30 Apr 2008 10:16:00 GMT

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

---

Hi,

maratrus wrote on Wed, 30 April 2008 10:07Hi,

I've suspected that something is not quite clear with ubuntu "ifdown\*" utilities.

But the fact is that in redhat based OS ifdown is a simple script so it's easy to investigate it: it calls appropriate script in /etc/sysconfig/network-scripts/.

And the only thing I've found there is:

Yep, that's true.

Quote:1. If it's possible we can investigate src files of appropriate utilities in ubuntu OS.

bridge : located in /etc/network/if-post-down.d/bridge

vlan : located : /etc/network/if-post-down.d/vlan

ifdown-venet : located : /etc/sysconfig/network-scripts/ifdown-venet

in /etc/sysconfig/network-scripts/ there are 2 others scripts : ifcfg-venet0 and ifup-venet and in /etc/network/if-down.d/ there is only postfix script but there are not used in our case i guess.

I've try ifdown without those script (move them to /tmp), but nothing change

Quote:2. ifdown can generate another process so could you possible try to it in the following way:

```
strace -fF -o out.txt ifdown eth1.35
```

see out.txt \*with\* script in /etc/network/if-post-down.d/ and /etc/sysconfig/network-scripts/

see out2.txt \*without\* script in /etc/network/if-post-down.d/ and /etc/sysconfig/network-scripts/

Thank you for your help

Yoann

## File Attachments

---

- 1) [out.txt](#), downloaded 543 times
  - 2) [bridge](#), downloaded 416 times
  - 3) [vlan](#), downloaded 493 times
  - 4) [ifdown-venet](#), downloaded 372 times
  - 5) [out2.txt](#), downloaded 500 times
- 

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs

Posted by [mistur](#) on Wed, 30 Apr 2008 11:05:26 GMT

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

---

Others strace of network command that are giving the same behavior :

```
strace -fF -o out2.txt ip link set eth1.35 down
```

```
23822 execve("/bin/ip", ["ip", "link", "set", "eth1.35", "down"], [/* 14 vars */]) = 0
23822 uname({sys="Linux", node="http03", ...}) = 0
23822 brk(0) = 0x523000
23822 mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS,
-1, 0) = 0x2ab089670000
23822 access("/etc/ld.so.nohwcap", F_OK) = -1 ENOENT (No such file or directory)
23822 mmap(NULL, 8192, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS,
-1, 0) = 0x2ab089671000
23822 access("/etc/ld.so.preload", R_OK) = -1 ENOENT (No such file or directory)
23822 open("/etc/ld.so.cache", O_RDONLY) = 3
23822 fstat(3, {st_mode=S_IFREG|0644, st_size=13937, ...}) = 0
23822 mmap(NULL, 13937, PROT_READ, MAP_PRIVATE, 3, 0) = 0x2ab089673000
23822 close(3) = 0
23822 access("/etc/ld.so.nohwcap", F_OK) = -1 ENOENT (No such file or directory)
23822 open("/lib/libresolv.so.2", O_RDONLY) = 3
23822 read(3, "\177ELF\2\1\1\0\0\0\0\0\0\0\0\0\0\3\0>\0\1\0\0\0\3406\0\0"... , 640) = 640
23822 fstat(3, {st_mode=S_IFREG|0644, st_size=75288, ...}) = 0
23822 mmap(NULL, 1131976, PROT_READ|PROT_EXEC, MAP_PRIVATE|MAP_DENYWRITE,
```

```

3, 0) = 0x2ab089771000
23822 mprotect(0x2ab089782000, 1062344, PROT_NONE) = 0
23822 mmap(0x2ab089882000, 4096, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_DENYWRITE, 3, 0x11000) = 0x2ab089882000
23822 mmap(0x2ab089883000, 9672, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_ANONYMOUS, -1, 0) = 0x2ab089883000
23822 close(3) = 0
23822 access("/etc/ld.so.nohwcap", F_OK) = -1 ENOENT (No such file or directory)
23822 open("/lib/libc.so.6", O_RDONLY) = 3
23822 read(3, "\177ELF\2\1\1\0\0\0\0\0\0\0\0\3\0>\0\1\0\0\0@\305\1\0"..., 640) = 640
23822 fstat(3, {st_mode=S_IFREG|0755, st_size=1267544, ...}) = 0
23822 mmap(NULL, 2327048, PROT_READ|PROT_EXEC, MAP_PRIVATE|MAP_DENYWRITE,
3, 0) = 0x2ab089886000
23822 mprotect(0x2ab0899a3000, 1159688, PROT_NONE) = 0
23822 mmap(0x2ab089aa3000, 94208, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_DENYWRITE, 3, 0x11d000) = 0x2ab089aa3000
23822 mmap(0x2ab089aba000, 16904, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_ANONYMOUS, -1, 0) = 0x2ab089aba000
23822 close(3) = 0
23822 mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS,
-1, 0) = 0x2ab089abf000
23822 arch_prctl(ARCH_SET_FS, 0x2ab089abf6d0) = 0
23822 munmap(0x2ab089673000, 13937) = 0
23822 socket(PF_INET, SOCK_DGRAM, IPPROTO_IP) = 3
23822 ioctl(3, SIOCGIFFLAGS, {ifr_name="eth1.35",
ifr_flags=IFF_UP|IFF_BROADCAST|IFF_RUNNING|IFF_MULTICAST}) = 0
23822 ioctl(3, SIOCSIFFLAGS, 0x7fff2144eef0) = 0
23822 close(3) = 0
23822 exit_group(0) = ?

```

```
strace -fF -o out2.txt ifconfig eth1.35 down
```

```

23877 execve("/sbin/ifconfig", ["ifconfig", "eth1.35", "down"], [/* 14 vars */]) = 0
23877 uname({sys="Linux", node="http03", ...}) = 0
23877 brk(0) = 0x511000
23877 mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS,
-1, 0) = 0x2b8ba95d7000
23877 access("/etc/ld.so.nohwcap", F_OK) = -1 ENOENT (No such file or directory)
23877 mmap(NULL, 8192, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS,
-1, 0) = 0x2b8ba95d8000
23877 access("/etc/ld.so.preload", R_OK) = -1 ENOENT (No such file or directory)
23877 open("/etc/ld.so.cache", O_RDONLY) = 3
23877 fstat(3, {st_mode=S_IFREG|0644, st_size=13937, ...}) = 0
23877 mmap(NULL, 13937, PROT_READ, MAP_PRIVATE, 3, 0) = 0x2b8ba95da000
23877 close(3) = 0
23877 access("/etc/ld.so.nohwcap", F_OK) = -1 ENOENT (No such file or directory)
23877 open("/lib/libc.so.6", O_RDONLY) = 3

```

```
23877 read(3, "\177ELF\2\1\1\0\0\0\0\0\0\0\0\0\0\03\0>\0\1\0\0\0@\305\1\0"..., 640) = 640
23877 fstat(3, {st_mode=S_IFREG|0755, st_size=1267544, ...}) = 0
23877 mmap(NULL, 2327048, PROT_READ|PROT_EXEC, MAP_PRIVATE|MAP_DENYWRITE,
3, 0) = 0x2b8ba96d8000
23877 mprotect(0x2b8ba97f5000, 1159688, PROT_NONE) = 0
23877 mmap(0x2b8ba98f5000, 94208, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_DENYWRITE, 3, 0x11d000) = 0x2b8ba98f5000
23877 mmap(0x2b8ba990c000, 16904, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_ANONYMOUS, -1, 0) = 0x2b8ba990c000
23877 close(3) = 0
23877 mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS,
-1, 0) = 0x2b8ba9911000
23877 arch_prctl(ARCH_SET_FS, 0x2b8ba99116d0) = 0
23877 munmap(0x2b8ba95da000, 13937) = 0
23877 open("/usr/lib/locale/locale-archive", O_RDONLY) = -1 ENOENT (No such file or directory)
23877 brk(0) = 0x511000
23877 brk(0x532000) = 0x532000
23877 open("/usr/share/locale/locale.alias", O_RDONLY) = 3
23877 fstat(3, {st_mode=S_IFREG|0644, st_size=2586, ...}) = 0
23877 mmap(NULL, 32768, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS,
-1, 0) = 0x2b8ba95da000
23877 read(3, "# Locale name alias data base.\n#"..., 32768) = 2586
23877 read(3, "", 32768) = 0
23877 close(3) = 0
23877 munmap(0x2b8ba95da000, 32768) = 0
23877 open("/usr/lib/locale/fr_FR@euro/LC_IDENTIFICATION", O_RDONLY) = 3
23877 fstat(3, {st_mode=S_IFREG|0644, st_size=444, ...}) = 0
23877 mmap(NULL, 444, PROT_READ, MAP_PRIVATE, 3, 0) = 0x2b8ba95da000
23877 close(3) = 0
23877 open("/usr/lib/locale/fr_FR@euro/LC_MEASUREMENT", O_RDONLY) = 3
23877 fstat(3, {st_mode=S_IFREG|0644, st_size=29, ...}) = 0
23877 mmap(NULL, 29, PROT_READ, MAP_PRIVATE, 3, 0) = 0x2b8ba95db000
23877 close(3) = 0
23877 open("/usr/lib/locale/fr_FR@euro/LC_TELEPHONE", O_RDONLY) = 3
23877 fstat(3, {st_mode=S_IFREG|0644, st_size=62, ...}) = 0
23877 mmap(NULL, 62, PROT_READ, MAP_PRIVATE, 3, 0) = 0x2b8ba95dc000
23877 close(3) = 0
23877 open("/usr/lib/locale/fr_FR@euro/LC_ADDRESS", O_RDONLY) = 3
23877 fstat(3, {st_mode=S_IFREG|0644, st_size=162, ...}) = 0
23877 mmap(NULL, 162, PROT_READ, MAP_PRIVATE, 3, 0) = 0x2b8ba95dd000
23877 close(3) = 0
23877 open("/usr/lib/locale/fr_FR@euro/LC_NAME", O_RDONLY) = 3
23877 fstat(3, {st_mode=S_IFREG|0644, st_size=77, ...}) = 0
23877 mmap(NULL, 77, PROT_READ, MAP_PRIVATE, 3, 0) = 0x2b8ba95de000
23877 close(3) = 0
23877 open("/usr/lib/locale/fr_FR@euro/LC_PAPER", O_RDONLY) = 3
23877 fstat(3, {st_mode=S_IFREG|0644, st_size=40, ...}) = 0
23877 mmap(NULL, 40, PROT_READ, MAP_PRIVATE, 3, 0) = 0x2b8ba95df000
```

```

23877 close(3) = 0
23877 open("/usr/lib/locale/fr_FR@euro/LC_MESSAGES", O_RDONLY) = 3
23877 fstat(3, {st_mode=S_IFDIR|0755, st_size=28, ...}) = 0
23877 close(3) = 0
23877 open("/usr/lib/locale/fr_FR@euro/LC_MESSAGES/SYS_LC_MESSAGES", O_RDONLY) =
3
23877 fstat(3, {st_mode=S_IFREG|0644, st_size=60, ...}) = 0
23877 mmap(NULL, 60, PROT_READ, MAP_PRIVATE, 3, 0) = 0x2b8ba95e0000
23877 close(3) = 0
23877 open("/usr/lib/locale/fr_FR@euro/LC_MONETARY", O_RDONLY) = 3
23877 fstat(3, {st_mode=S_IFREG|0644, st_size=292, ...}) = 0
23877 mmap(NULL, 292, PROT_READ, MAP_PRIVATE, 3, 0) = 0x2b8ba95e1000
23877 close(3) = 0
23877 open("/usr/lib/locale/fr_FR@euro/LC_COLLATE", O_RDONLY) = 3
23877 fstat(3, {st_mode=S_IFREG|0644, st_size=20544, ...}) = 0
23877 mmap(NULL, 20544, PROT_READ, MAP_PRIVATE, 3, 0) = 0x2b8ba95e2000
23877 close(3) = 0
23877 open("/usr/lib/locale/fr_FR@euro/LC_TIME", O_RDONLY) = 3
23877 fstat(3, {st_mode=S_IFREG|0644, st_size=2388, ...}) = 0
23877 mmap(NULL, 2388, PROT_READ, MAP_PRIVATE, 3, 0) = 0x2b8ba95e8000
23877 close(3) = 0
23877 open("/usr/lib/locale/fr_FR@euro/LC_NUMERIC", O_RDONLY) = 3
23877 fstat(3, {st_mode=S_IFREG|0644, st_size=60, ...}) = 0
23877 mmap(NULL, 60, PROT_READ, MAP_PRIVATE, 3, 0) = 0x2b8ba95e9000
23877 close(3) = 0
23877 open("/usr/lib/locale/fr_FR@euro/LC_CTYPE", O_RDONLY) = 3
23877 fstat(3, {st_mode=S_IFREG|0644, st_size=207972, ...}) = 0
23877 mmap(NULL, 207972, PROT_READ, MAP_PRIVATE, 3, 0) = 0x2b8ba95ea000
23877 close(3) = 0
23877 uname({sys="Linux", node="http03", ...}) = 0
23877 access("/proc/net", R_OK) = 0
23877 access("/proc/net/unix", R_OK) = 0
23877 socket(PF_FILE, SOCK_DGRAM, 0) = 3
23877 socket(PF_INET, SOCK_DGRAM, IPPROTO_IP) = 4
23877 access("/proc/net/af_inet6", R_OK) = 0
23877 socket(PF_INET6, SOCK_DGRAM, IPPROTO_IP) = 5
23877 access("/proc/net/ax25", R_OK) = -1 ENOENT (No such file or directory)
23877 access("/proc/net/nr", R_OK) = -1 ENOENT (No such file or directory)
23877 access("/proc/net/rose", R_OK) = -1 ENOENT (No such file or directory)
23877 access("/proc/net/ipx", R_OK) = -1 ENOENT (No such file or directory)
23877 access("/proc/net/appletalk", R_OK) = -1 ENOENT (No such file or directory)
23877 access("/proc/sys/net/econet", R_OK) = -1 ENOENT (No such file or directory)
23877 access("/proc/sys/net/ash", R_OK) = -1 ENOENT (No such file or directory)
23877 access("/proc/net/x25", R_OK) = -1 ENOENT (No such file or directory)
23877 open("/usr/share/locale/fr_FR@euro/LC_MESSAGES/net-tools.mo", O_RDONLY) = -1
ENOENT (No such file or directory)
23877 open("/usr/share/locale/fr@euro/LC_MESSAGES/net-tools.mo", O_RDONLY) = -1
ENOENT (No such file or directory)

```

```
23877 open("/usr/share/locale/fr_FR/LC_MESSAGES/net-tools.mo", O_RDONLY) = -1 ENOENT
(No such file or directory)
23877 open("/usr/share/locale/fr/LC_MESSAGES/net-tools.mo", O_RDONLY) = -1 ENOENT (No
such file or directory)
23877 open("/usr/share/locale-langpack/fr_FR@euro/LC_MESSAGES/net-tools.mo",
O_RDONLY) = -1 ENOENT (No such file or directory)
23877 open("/usr/share/locale-langpack/fr@euro/LC_MESSAGES/net-tools.mo", O_RDONLY) =
-1 ENOENT (No such file or directory)
23877 open("/usr/share/locale-langpack/fr_FR/LC_MESSAGES/net-tools.mo", O_RDONLY) = -1
ENOENT (No such file or directory)
23877 open("/usr/share/locale-langpack/fr/LC_MESSAGES/net-tools.mo", O_RDONLY) = -1
ENOENT (No such file or directory)
23877 ioctl(4, SIOCGIFFLAGS, {ifr_name="eth1.35",
ifr_flags=IFF_UP|IFF_BROADCAST|IFF_RUNNING|IFF_MULTICAST}) = 0
23877 ioctl(4, SIOCSIFFLAGS, 0x7fff014e7f60) = 0
23877 exit_group(0)                = ?
```

---

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs

Posted by [maratrus](#) on Wed, 30 Apr 2008 11:15:54 GMT

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

---

Hi,

1. But we've get rid of the message "unregister\_vlan\_device: ERROR: Tried to remove a non-vlan device with VLAN code, name: eth1.35 priv\_flags: 0", haven't we?

out.txt

Quote:

```
execve("/sbin/vconfig", ["vconfig", "rem", "eth1.35"], [/* 11 vars */]) = 0
```

out2.txt

Quote:

```
execve("/sbin/ifconfig", ["ifconfig", "eth1.35", "inet", "0"], [/* 15 vars */]) = 0
```

So, in the first case we call vconfig to remove non-vlan device with a VLAN code; in the second case we call ifconfig to put down the interface.

2. When you run "ifconfig eth1.35 down" inside the ubuntu VPS other VPSes loose an access to nfs share. Is that true? And when you run "ifconfig eth1.32 down"?

3. Anything in logs or in dmesg?

P.S. I don't expect this help but could you possible rename your eth1.35 interface inside VE to eth1?

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs

Posted by [mistur](#) on Wed, 30 Apr 2008 13:40:23 GMT

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

---

maratrus wrote on Wed, 30 April 2008 13:151. But we've get rid of the message "unregister\_vlan\_device: ERROR: Tried to remove a non-vlan device with VLAN code, name: eth1.35 priv\_flags: 0", haven't we?

in fact, inside a VE eth1.35 is not eth1 with vlan 35 tag, it is viewed as a physical interface for the OS, so /proc/net/vlan/config is empty, if I understand well.

Quote:So, in the first case we call vconfig to remove non-vlan device with a VLAN code; in the second case we call ifconfig to put down the interface.

ubuntu script ckeck if the interface is associate to a vlan by the name : ethX.Y where Y is the number of the vlan

but in this case, the number after the interface does not refer to the vlan conf for the VE but refere to the bridge associate to the vlan in the HN.

Quote:2. When you run "ifconfig eth1.35 down" inside the ubuntu VPS other VPSes loose an access to nfs share. Is that true? And when you run "ifconfig eth1.32 down"?

I'm loosing only one nfs share out of the 3 used, I can ping the nfs server, I can ssh the nfs server, I can read or write on the 2 others nfs share, but one specifically die...

that's the really strange point...

Quote:3. Anything in logs or in dmesg?

I got this in the dmesg on the HN

```
[258012.479557] unregister_vlan_device: ERROR: Tried to remove a non-vlan device with VLAN code, name: eth1.35 priv_flags: 0
```

```
[258012.646468] unregister_vlan_device: ERROR: Tried to remove a non-vlan device with VLAN code, name: eth1.32 priv_flags: 0
```

```
[258013.466439] nfs: RPC call returned error 101
```

```
[258013.711497] nfs: RPC call returned error 101
```

```
[258013.715907] nfs: RPC call returned error 101
```

Quote:P.S. I don't expect this help but could you possible rename your eth1.35 interface inside VE to eth1?

I can't due to the vlan configuration on switches, but I've add just eth1 to the VE which not exist before (only eth1.35 and eth1.32), but nothing has changed.

all tests are done on a test server but over a network setup where is the production

so, I can really use vlan inside VE for now.

Yoann, still investigating

---

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs

Posted by [maratrus](#) on Wed, 30 Apr 2008 18:06:24 GMT

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

---

Hi,

I haven't understood you clearly:

Quote:

in fact, inside a VE eth1.35 is not eth1 with vlan 35 tag, it is viewed as a physical interface for the OS, so /proc/net/vlan/config is empty, if I understand well.

Yes, I think in the same way. According to your configuration.

Quote:

ubuntu script ckeck if the interface is associate to a vlan by the name : ethX.Y where Y is the number of the vlan

but in this case, the number after the interface does not refer to the vlan conf for the VE but refere to the bridge associate to the vlan in the HN.

So, I think we are talking about vlan script. Yes, I agree with you. It checks some conditions and then calls "vconfig" utility. I haven't looked at "vconfig" sources. May be it does another checks, I don't know but I have a question:

Have you got this message

Quote:

```
unregister_vlan_device: ERROR: Tried to remove a non-vlan device with VLAN code, name:
eth1.35 priv_flags: 0
```

in the both cases (before and after you move /etc/network/if-post-down.d/ scripts)

You've mentioned that "ifconfig eth1.35 down" reproduces the issue. Could you possibly say does "ifconfig eth1.32" reproduce the issue?

Quote:

I can't due to the vlan configuration on switches

I'm not clearly understand why can't you use "--netif\_add eth1" instead of "--netif\_add eth1.35"

P.S. I'll ask another people. I hope somebody could help us with this problem.

---

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs

Posted by [mistur](#) on Wed, 30 Apr 2008 20:10:15 GMT

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

---

Quote:Quote:ubuntu script ckeck if the interface is associate to a vlan by the name : ethX.Y where Y is the number of the vlan

but in this case, the number after the interface does not refer to the vlan conf for the VE but refere to the bridge associate to the vlan in the HN.

So, I think we are talking about vlan script. Yes, I agree with you. It checks some conditions and then calls "vconfig" utility. I haven't looked at "vconfig" sources. May be it does another checks, I don't know but I have a question:

Have you got this message

Quote:unregister\_vlan\_device: ERROR: Tried to remove a non-vlan device with VLAN code, name: eth1.35 priv\_flags: 0

in the both cases (before and after you move /etc/network/if-post-down.d/ scripts)

You've mentioned that "ifconfig eth1.35 down" reproduces the issue. Could you possibly say does "ifconfig eth1.32" reproduce the issue?

nop, I can't reproduce the issue with ifconfig eth1.32 down.

Quote:Quote:I can't due to the vlan configuration on switches

I'm not clearly understand why can't you use "--netif\_add eth1" instead of "--netif\_add eth1.35"

Ok, I understand. if I change eth1.35 by eth1, I can't reproduce the issue.

so I've juste remove the reference of the vlan in the name of the interface inside the VE, but nothing has been changed on HN side. eth1 is still associated to veth204.35, and on the HN, veth204.35 and eth1.35 are in the bridge brvlan35.

Quote:P.S. I'll ask another people. I hope somebody could help us with this problem.

thanks

Yoann

---

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs  
Posted by [maratrus](#) on Sun, 04 May 2008 11:15:39 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hi,

I asked about your problem and the answer was the following:

1. Disable

```
CONFIG_NFS_V4,  
CONFIG_NFSD  
CONFIG_NFSD_V2_ACL  
CONFIG_NFSD_V3  
CONFIG_NFSD_V3_ACL  
CONFIG_NFSD_V4  
CONFIG_NFSD_TCP
```

and recompile the kernel, please?

2. Am I right that your NFS servers are running on the external node?

---

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs  
Posted by [mistur](#) on Sun, 04 May 2008 16:27:39 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

I will do it next week and give you feed back asap.

thx

Yoann

---

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs  
Posted by [mistur](#) on Tue, 13 May 2008 14:24:05 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

kernel recompiled but problem still present

```
root@vz01:/boot# uname -a  
Linux vz01 2.6.18-028stab053 #1 SMP Tue May 13 13:47:36 CEST 2008 x86_64 GNU/Linux  
root@vz01:/boot# grep -i nfs config-2.6.18-028stab053  
CONFIG_NFS_FS=y  
CONFIG_NFS_V3=y  
CONFIG_NFS_V3_ACL=y  
# CONFIG_NFS_V4 is not set  
CONFIG_NFS_DIRECTIO=y  
# CONFIG_NFSD is not set
```

CONFIG\_ROOT\_NFS=y  
CONFIG\_NFS\_ACL\_SUPPORT=y  
CONFIG\_NFS\_COMMON=y  
CONFIG\_NCPFS\_NFS\_NS=y

Yoann

---

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs  
Posted by [mistur](#) on Tue, 20 May 2008 16:35:11 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

I don't have new idea that can help to find where does this problems come from.

is anyone can help me on this problem ?

Yoann

---

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs  
Posted by [maratrus](#) on Mon, 26 May 2008 13:45:54 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hello,

1. Please show how do you exactly create eth\* interfaces per VE. (the exact command - "vzctl set --netif\_add ...")
  2. And what is the order of interfaces you add to the bridge?  
(for example: the first is eth1.35@eth1, then veth101.35, veth102.35 etc) And which of them cause the problem? I mean if the problem VE was the first to add?
  3. Let's assume you've started several VEs and mounted NFS partition inside each of them. Please show "ip a l", "ip r l", "brctl show" from HN at that moment. Then you can break this configuration via "ifconfig eth1.35 down" from inside the problem VE. Please, do it. And then show the previous commands from HN.
  4. Let's assume you've broken the configuration. Do other VEs keep a network connectivity?
- 

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs  
Posted by [mistur](#) on Mon, 26 May 2008 16:40:22 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

maratrus wrote on Mon, 26 May 2008 15:45Hello,

Hi,

maratrus wrote on Mon, 26 May 2008 15:45. Please show how do you exactly create eth\* interfaces per VE. (the exact command - "vzctl set --netif\_add ...")

NETIF in all XXX.conf where XXX = VEID

```
NETIF="ifname=eth1.32,mac=00:16:3E:0X:XX:1E,host_ifname=vethXXX.32,host_mac=FE:FF:FF:FF:FF:FF;ifname=eth1.35,mac=00:16:3E:0X:XX:1E,host_ifname=vethXXX.35,host_mac=FE:FF:FF:FF:FF:FF;ifname=eth1.36,mac=00:16:3E:0X:XX:1E,host_ifname=vethXXX.36,host_mac=FE:FF:FF:FF:FF:FF"
```

/etc/vz/vznet.conf :

```
#!/bin/bash
EXTERNAL_SCRIPT="/usr/sbin/vznetaddroute
```

here the script that create the interface inside the HN

```
#!/bin/bash
# /usr/sbin/vznetaddroute
# a script to bring up virtual network interfaces (veth's) in a VE
CONFIGFILE=/etc/vz/conf/$VEID.conf
. $CONFIGFILE
VETH=$3
VZHOSTIF=`echo ${VETH} | cut -f1 -d"."`
VLAN=`echo ${VETH} | cut -f2 -sd"."`

if [ -n "${VLAN}" ] ; then
    VETH_BRIDGE=brvlan${VLAN}

    if [ ! -n "${VZHOSTIF}.${VLAN}" ] ; then
        echo "According to $CONFIGFILE VE$VEID has no veth interface configured."
        exit 1
    fi

    if [ ! -n "${VETH_BRIDGE}" ] ; then
        echo "According to $CONFIGFILE VE$VEID has no veth VETH_BRIDGE configured."
        exit 1
    fi

    echo "Adding interface ${VZHOSTIF}.${VLAN} to VETH_BRIDGE ${VETH_BRIDGE} on VE0"
```

```

if [ -z "/usr/sbin/brctl show | /bin/grep ${VETH_BRIDGE}" ]; then
    /usr/sbin/brctl addbr ${VETH_BRIDGE}
    /usr/sbin/brctl addif ${VETH_BRIDGE} eth1.${VLAN}
    /sbin/ip link set ${VETH_BRIDGE} up
fi

/usr/sbin/brctl addif ${VETH_BRIDGE} ${VZHOSTIF}.${VLAN}
/sbin/ip link set ${VZHOSTIF}.${VLAN} up
else
    VETH_BRIDGE=brvlan1

if [ ! -n "${VZHOSTIF}" ]; then
    echo "According to $CONFIGFILE VE$VEID has no veth interface configured."
    exit 1
fi

if [ ! -n "${VETH_BRIDGE}" ]; then
    echo "According to $CONFIGFILE VE$VEID has no veth VETH_BRIDGE configured."
    exit 1
fi

echo "Adding interface ${VZHOSTIF} to VETH_BRIDGE ${VETH_BRIDGE} on VE0"

/usr/sbin/brctl addif ${VETH_BRIDGE} ${VZHOSTIF}
/sbin/ip link set ${VZHOSTIF} up

fi

exit 0

```

maratrus wrote on Mon, 26 May 2008 15:452. And what is the order of interfaces you add to the bridge?

(for example: the first is eth1.35@eth1, then veth101.35, veth102.35 etc) And which of them cause the problem? I mean if the problem VE was the first to add?

each VE is started one by one, interface are mounted in the ordre the are appear in the NETIF parameter in the conf file.

the if which create problems is the second one.

maratrus wrote on Mon, 26 May 2008 15:453. Let's assume you've started several VEs and mounted NFS partition inside each of them. Please show "ip a l", "ip r l", "brctl show" from HN at that moment.

Then you can break this configuration via "ifconfig eth1.35 down" from inside the problem VE. Please, do it. And then show the previous commands from HN.

I can reproduce the bug with only 2 VEs up, so I guess it's enough

With all VEs working well :

```
# ip a l
2: lo: <LOOPBACK,UP,10000> mtu 16436 qdisc noqueue
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
4: eth0: <BROADCAST,MULTICAST> mtu 1500 qdisc noop qlen 1000
    link/ether 00:e0:81:4b:2f:da brd ff:ff:ff:ff:ff:ff
6: eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
8: eth1.35@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
10: brvlan35: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet 10.0.0.247/24 brd 10.0.0.255 scope global brvlan35
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
12: eth1.15@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
14: brvlan15: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
16: eth1.17@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
18: brvlan17: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
20: eth1.19@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
22: brvlan19: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
```

```
valid_lft forever preferred_lft forever
24: eth1.20@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
valid_lft forever preferred_lft forever
26: brvlan20: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
valid_lft forever preferred_lft forever
28: eth1.30@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
valid_lft forever preferred_lft forever
30: brvlan30: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
valid_lft forever preferred_lft forever
32: eth1.32@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
valid_lft forever preferred_lft forever
34: brvlan32: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
valid_lft forever preferred_lft forever
36: eth1.36@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
valid_lft forever preferred_lft forever
38: brvlan36: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
valid_lft forever preferred_lft forever
40: eth1.39@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
valid_lft forever preferred_lft forever
42: brvlan39: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
inet 192.168.25.180/28 brd 192.168.25.191 scope global brvlan39
inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
valid_lft forever preferred_lft forever
44: sit0: <NOARP> mtu 1480 qdisc noop
link/sit 0.0.0.0 brd 0.0.0.0
46: eth1.218@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
valid_lft forever preferred_lft forever
```

```

48: brvlan218: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff
  inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
    valid_lft forever preferred_lft forever
1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,10000> mtu 1500 qdisc noqueue
  link/void
11: veth203.32: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether fe:ff:ff:ff:ff:ff brd ff:ff:ff:ff:ff
  inet6 fe80::fcff:ffff:feff:ffff/64 scope link
    valid_lft forever preferred_lft forever
13: veth203.35: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether fe:ff:ff:ff:ff:ff brd ff:ff:ff:ff:ff
  inet6 fe80::fcff:ffff:feff:ffff/64 scope link
    valid_lft forever preferred_lft forever
15: veth203.36: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether fe:ff:ff:ff:ff:ff brd ff:ff:ff:ff:ff
  inet6 fe80::fcff:ffff:feff:ffff/64 scope link
    valid_lft forever preferred_lft forever
17: veth204.32: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether fe:ff:ff:ff:ff:ff brd ff:ff:ff:ff:ff
  inet6 fe80::fcff:ffff:feff:ffff/64 scope link
    valid_lft forever preferred_lft forever
19: veth204.35: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether fe:ff:ff:ff:ff:ff brd ff:ff:ff:ff:ff
  inet6 fe80::fcff:ffff:feff:ffff/64 scope link
    valid_lft forever preferred_lft forever
21: veth204.36: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether fe:ff:ff:ff:ff:ff brd ff:ff:ff:ff:ff
  inet6 fe80::fcff:ffff:feff:ffff/64 scope link
    valid_lft forever preferred_lft forever

```

```
# ip r l
```

```

192.168.25.176/28 dev brvlan39 proto kernel scope link src 192.168.25.180
10.0.0.0/24 dev brvlan35 proto kernel scope link src 10.0.0.247
default via 192.168.25.177 dev brvlan39

```

```
# brctl show
```

bridge name	bridge id	STP enabled	interfaces
brvlan35	8000.00e0814b2fdb	no	eth1.35 veth203.35 veth204.35
brvlan15	8000.00e0814b2fdb	no	eth1.15
brvlan17	8000.00e0814b2fdb	no	eth1.17
brvlan19	8000.00e0814b2fdb	no	eth1.19
brvlan20	8000.00e0814b2fdb	no	eth1.20
brvlan30	8000.00e0814b2fdb	no	eth1.30
brvlan32	8000.00e0814b2fdb	no	eth1.32 veth203.32

		veth204.32	
brvlan36	8000.00e0814b2fdb	no	eth1.36
		veth203.36	
		veth204.36	
brvlan39	8000.00e0814b2fdb	no	eth1.39
brvlan218	8000.00e0814b2fdb	no	eth1.218

after an ifconfig eth1.35 down in one of the VEs with the nfs mount point borked

```
# ip a l ;
2: lo: <LOOPBACK,UP,10000> mtu 16436 qdisc noqueue
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
4: eth0: <BROADCAST,MULTICAST> mtu 1500 qdisc noop qlen 1000
    link/ether 00:e0:81:4b:2f:da brd ff:ff:ff:ff:ff:ff
6: eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
8: eth1.35@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
10: brvlan35: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet 10.0.0.247/24 brd 10.0.0.255 scope global brvlan35
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
12: eth1.15@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
14: brvlan15: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
16: eth1.17@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
18: brvlan17: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
```

```

inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
  valid_lft forever preferred_lft forever
20: eth1.19@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
  inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
    valid_lft forever preferred_lft forever
22: brvlan19: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
  inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
    valid_lft forever preferred_lft forever
24: eth1.20@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
  inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
    valid_lft forever preferred_lft forever
26: brvlan20: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
  inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
    valid_lft forever preferred_lft forever
28: eth1.30@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
  inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
    valid_lft forever preferred_lft forever
30: brvlan30: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
  inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
    valid_lft forever preferred_lft forever
32: eth1.32@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
  inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
    valid_lft forever preferred_lft forever
34: brvlan32: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
  inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
    valid_lft forever preferred_lft forever
36: eth1.36@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
  inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
    valid_lft forever preferred_lft forever
38: brvlan36: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
  inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
    valid_lft forever preferred_lft forever
40: eth1.39@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
  inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
    valid_lft forever preferred_lft forever
42: brvlan39: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
  link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff

```

```

inet 192.168.25.180/28 brd 192.168.25.191 scope global brvlan39
inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
    valid_lft forever preferred_lft forever
44: sit0: <NOARP> mtu 1480 qdisc noop
    link/sit 0.0.0.0 brd 0.0.0.0
46: eth1.218@eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
48: brvlan218: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:e0:81:4b:2f:db brd ff:ff:ff:ff:ff:ff
    inet6 fe80::2e0:81ff:fe4b:2fdb/64 scope link
        valid_lft forever preferred_lft forever
1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,10000> mtu 1500 qdisc noqueue
    link/void
11: veth203.32: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether fe:ff:ff:ff:ff:ff brd ff:ff:ff:ff:ff:ff
    inet6 fe80::fcff:ffff:feff:ffff/64 scope link
        valid_lft forever preferred_lft forever
13: veth203.35: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether fe:ff:ff:ff:ff:ff brd ff:ff:ff:ff:ff:ff
    inet6 fe80::fcff:ffff:feff:ffff/64 scope link
        valid_lft forever preferred_lft forever
15: veth203.36: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether fe:ff:ff:ff:ff:ff brd ff:ff:ff:ff:ff:ff
    inet6 fe80::fcff:ffff:feff:ffff/64 scope link
        valid_lft forever preferred_lft forever
17: veth204.32: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether fe:ff:ff:ff:ff:ff brd ff:ff:ff:ff:ff:ff
    inet6 fe80::fcff:ffff:feff:ffff/64 scope link
        valid_lft forever preferred_lft forever
19: veth204.35: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether fe:ff:ff:ff:ff:ff brd ff:ff:ff:ff:ff:ff
    inet6 fe80::fcff:ffff:feff:ffff/64 scope link
        valid_lft forever preferred_lft forever
21: veth204.36: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether fe:ff:ff:ff:ff:ff brd ff:ff:ff:ff:ff:ff
    inet6 fe80::fcff:ffff:feff:ffff/64 scope link
        valid_lft forever preferred_lft forever

```

```
# ip r l
```

```

192.168.25.176/28 dev brvlan39 proto kernel scope link src 192.168.25.180
10.0.0.0/24 dev brvlan35 proto kernel scope link src 10.0.0.247
default via 192.168.25.177 dev brvlan39

```

```
# brctl show
```

bridge name	bridge id	STP enabled	interfaces
brvlan35	8000.00e0814b2fdb	no	eth1.35

```

                                veth203.35
                                veth204.35
brvlan15      8000.00e0814b2fdb  no      eth1.15
brvlan17      8000.00e0814b2fdb  no      eth1.17
brvlan19      8000.00e0814b2fdb  no      eth1.19
brvlan20      8000.00e0814b2fdb  no      eth1.20
brvlan30      8000.00e0814b2fdb  no      eth1.30
brvlan32      8000.00e0814b2fdb  no      eth1.32
                                veth203.32
                                veth204.32
brvlan36      8000.00e0814b2fdb  no      eth1.36
                                veth203.36
                                veth204.36
brvlan39      8000.00e0814b2fdb  no      eth1.39
brvlan218     8000.00e0814b2fdb  no      eth1.218

```

maratrus wrote on Mon, 26 May 2008 15:454. Let's assume you've broken the configuration. Do other VEs keep a network connectivity?

well, infact others VEs does not loose their network connectivities, I can ping the nfs server, I can ssh it, other nfs mount point works well from the same servers, but with another IP

there are 8 ips on the nfs server dedicated respectivly for 8 nfs exports, the 8 ip are on the same vlan35.

they just loose acces to the one of the nfs partition, but not all of them...

for example I have

```

IP1:/mnt/data1 /data1
IP2:/mnt/data2 /data2
IP3:/mnt/data3 /data3

```

there is only one mount point broked (third for example) in my case,the 2 others are working well on others VE even if I perform a down on the vlan35 interface on one of the ubuntu based VE (except the VE where ifdown is performed of course).

BUT, I can remount the nfs export with another IP like IPX:/mnt/data3 /data3 where X is different to 3...

What I think is :

I might not be directly openvz related, but a side effect in nfs kernel land, like I have the same mount with exactly the same parameter (IP3:/mnt/data3 /data3) in each VEs.

when I performe a ifdown on eth1.35 in the VE, ubuntu script try to remove vlan35 for kernel, but

VE is not authorized due to all VE and the HN are using vlan35, and eth1.35 is not eth1 tagged with vlan 35 inside the VE it's the name of the interface like eth1 eth2, etc...

when the ubuntu scripts is trying to remove vlan35, in kernel land, there are a dead lock ore something similar.

when I umount all /data3, that's removing all data3 ref in the kernel, and lock is removed, then, I mount again data3

why other /dataX are not affected I don't know, maybe it related to the order the kernel keep mount point in memory.

I've try to change the order in fstab without success

```
IP3:/mnt/data3 /data3
IP1:/mnt/data1 /data1
IP2:/mnt/data2 /data2
```

hope that can help !

Yoann

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs

Posted by [mistur](#) on Mon, 26 May 2008 16:45:28 GMT

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

maybe there are a link with this issue for <http://forum.openvz.org/index.php?t=msg&th=6151&start=0> ?

regards

Yoann

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs

Posted by [maratrus](#) on Mon, 26 May 2008 17:17:51 GMT

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

Quote:

ubuntu script try to remove vlan35 for kernel,

So, does the command "ifconfig eth1.35 down" cause the message "unregister\_vlan\_device: ERROR: Tried to remove a non-vlan device with VLAN code, name: eth1.35 priv\_flags: 0" or this message doesn't appear.

You have shown strace output of that command

[http://forum.openvz.org/index.php?t=msg&&th=5949&goto=29539#msg\\_29932](http://forum.openvz.org/index.php?t=msg&&th=5949&goto=29539#msg_29932)

I haven't been able to find anything related to VLAN.

What's the difference between strace "ifconfig eth1.35 down" and "ifconfig eth1 down" (after eth1.35 rename to eth1)? Why the second case doesn't break the system? Where is the difference in strace output?

I think this sort of investigation could help us. It's the good for us to understand the difference between the "good" and the "bad" cases.

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs

Posted by [mistur](#) on Mon, 26 May 2008 17:43:10 GMT

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

---

maratrus wrote on Mon, 26 May 2008 19:17Quote:  
ubuntu script try to remove vlan35 for kernel,

So, does the command "ifconfig eth1.35 down" cause the message "unregister\_vlan\_device: ERROR: Tried to remove a non-vlan device with VLAN code, name: eth1.35 priv\_flags: 0" or this message doesn't appear.

You have shown strace output of that command

[http://forum.openvz.org/index.php?t=msg&&th=5949&goto=29539#msg\\_29932](http://forum.openvz.org/index.php?t=msg&&th=5949&goto=29539#msg_29932)

I haven't been able to find anything related to VLAN.

inside a VE :

```
# ifdown eth1.35
Internet Systems Consortium DHCP Client V3.0.3
Copyright 2004-2005 Internet Systems Consortium.
All rights reserved.
For info, please visit http://www.isc.org/products/DHCP
```

```
Listening on LPF/eth1.35/00:16:3e:02:03:1e
Sending on LPF/eth1.35/00:16:3e:02:03:1e
Sending on Socket/fallback
DHCPRELEASE on eth1.35 to 10.0.0.1 port 67
ERROR: trying to remove VLAN -:eth1.35:- error: Operation not permitted
# ifup eth1.35
Set name-type for VLAN subsystem. Should be visible in /proc/net/vlan/config
Device "eth1" does not exist.
eth1 does not exist, unable to create eth1.35
run-parts: /etc/network/if-pre-up.d/vlan exited with return code 1
Internet Systems Consortium DHCP Client V3.0.3
Copyright 2004-2005 Internet Systems Consortium.
```

All rights reserved.  
For info, please visit <http://www.isc.org/products/DHCP>

```
Listening on LPF/eth1.35/00:16:3e:02:03:1e
Sending on LPF/eth1.35/00:16:3e:02:03:1e
Sending on Socket/fallback
DHCPDISCOVER on eth1.35 to 255.255.255.255 port 67 interval 3
DHCPOFFER from 10.0.0.1
DHCPREQUEST on eth1.35 to 255.255.255.255 port 67
DHCPACK from 10.0.0.1
bound to 10.0.0.197 -- renewal in 3530590 seconds.
# ifconfig eth1.35 down
# ifconfig eth1.35 up
# ip l set down eth1.35
# ip l set up eth1.35
#
```

ifdown, ifconfig and ip l set down reproduce this issue, but "unregister\_vlan\_device: ERROR: ..." appear only with ifdown

"/etc/init.d/networking stop" uses ifdown

so maybe it's not vlan related, and I'm working on the wrong way.

maratrus wrote on Mon, 26 May 2008 19:17: What's the difference between strace "ifconfig eth1.35 down" and "ifconfig eth1 down" (after eth1.35 rename to eth1)? Why the second case doesn't break the system? Where is the difference in strace output?

If I had the response...

there are no big diff between the 2 straces... ifconfig down does not break directly nfs export I guess, but by shutdown the interface (whatever the way to set down), something is broken in kernel land. If not, I can't see how an action inside a VE can affect another VE.

maratrus wrote on Mon, 26 May 2008 19:17: I think this sort of investigation could help us. It's the good for us to understand the difference between the "good" and the "bad" cases.

true

Yoann

---

Subject: Re: Pb with multiple identical nfs mount dir in different VEs  
Posted by [mistur](#) on Tue, 27 May 2008 14:44:31 GMT  
[View Forum Message](#) <> [Reply to Message](#)

that patch has fixed the problem \o/.

[http://forum.openvz.org/index.php?t=msg&&th=6151&goto=30491#msg\\_30491](http://forum.openvz.org/index.php?t=msg&&th=6151&goto=30491#msg_30491)

thanks maratrus for your help and thanks Denis for the patch

Regards,

Yoann, happy

---