
Subject: SOLVED * Veth on startup (again) + bridge config - venet config

Posted by [sylvain](#) on Mon, 20 Nov 2006 04:31:43 GMT

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

Hi all,

I am using 5 VEs based on "centos-4-i386-minimal" template. I needed broadcast and I didn't care much about security so I configured my Ethernet device on each VE that I bridged on VE0: everything works fine!

Now I'd like to configure everything automatically when I run "vzctl start VEID".

Sounds like an old thread? Yes, another very similar question was asked. Installing the last version of vzctl would solve the problem. I already have the last version of the tool, however I still have a problem.

Here are a list of my problems:

venet0 is always configured inside the VE. I removed /etc/sysconfig/network-script/ifcfg-venet0 and it still appears at the next "vzctl start": I would like to get rid of this, and I would like to setup a proper routing for vethX inside the the VE.

it would be good if vethX could be added to vzbr0 when I run "vzctl start VEID". Is there a way to do this?

Here are the details of my configuration files:

A typical VEID.conf (without the copyright, thought)

```
ONBOOT="no"
```

```
# UBC parameters (in form of barrier:limit)
```

```
# Primary parameters
```

```
AVNUMPROC="40:40"
```

```
NUMPROC="65:65"
```

```
NUMTCP SOCK="80:80"
```

```
NUMOTHER SOCK="80:80"
```

```
VMGUAR PAGES="6144:2147483647"
```

```
# Secondary parameters
```

```
KMEMSIZE="2752512:2936012"
```

```
TCPSNDBUF="319488:524288"
```

```
TCPRCVBUF="319488:524288"
```

```
OTHERSOCKBUF="132096:336896"
```

```
DGRAMRCVBUF="132096:132096"
```

```
OOMGUAR PAGES="6144:2147483647"
```

```
# Auxiliary parameters
```

```
LOCKEDPAGES="32:32"
```

```
SHMPAGES="8192:8192"
PRIVVMPAGES="49152:53575"
NUMFILE="2048:2048"
NUMFLOCK="100:110"
NUMPTY="16:16"
NUMSIGINFO="256:256"
DCACHESIZE="1048576:1097728"
PHYSPPAGES="0:2147483647"
NUMIPTENT="128:128"
```

Disk quota parameters (in form of softlimit:hardlimit)

```
DISKSPACE="1048576:1153434"
DISKINODES="200000:220000"
QUOTATIME="0"
```

CPU fair sheduler parameter

```
CPUUNITS="1000"
VE_ROOT="/users/home/vz/root/$VEID"
VE_PRIVATE="/users/home/vz/private/$VEID"
OSTEMPLATE="centos-4-i386-minimal"
IP_ADDRESS=""
ORIGIN_SAMPLE="vps.basic"
HOSTNAME="tcp1"
VETH="veth1,00:12:34:56:78:1A,eth1,00:12:34:56:78:1B "
```

Inside the corresponding VE, /etc/sysconfig/network-script/ifcfg-eth1

```
DEVICE=eth1
BOOTPROTO=none
BROADCAST=192.168.7.255
HWADDR=00:12:34:56:78:1B
IPADDR=192.168.7.1
NETMASK=255.255.255.0
NETWORK=192.168.7.0
ONBOOT=yes
TYPE=Ethernet
USERCTL=no
PEERDNS=yes
IPV6INIT=no
GATEWAY=none
```

It is indeed parsed at startup of the VE.

However, since it also parse ifcfg-venet0, I end up with this routing:

Kernel IP routing table

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
-------------	---------	---------	-------	--------	-----	-----	-------

```
192.168.7.0 *          255.255.255.0 U  0  0    0 eth1
191.255.255.0 *       255.255.255.0 U  0  0    0 venet0
169.254.0.0 *         255.255.0.0   U  0  0    0 venet0
default    191.255.255.1 0.0.0.0      UG  0  0    0 venet0
```

Instead of this:

Kernel IP routing table

```
Destination Gateway Genmask Flags Metric Ref Use Iface
192.168.7.0 *          255.255.255.0 U  0  0    0 eth1
default *          0.0.0.0      UG  0  0    0 eth1
```

Also I would like to get rid of the line GATEWAY in /etc/sysconfig/networking inside the VE:

```
NETWORKING="yes"
HOSTNAME="tcp1"
GATEWAY="191.255.255.1"
```

Finally I don't know when to execute "brctl add vzbr0 veth1".

Any help is much appreciated . VE0 is a RHEL4

Subject: Re: Veth on startup (again) + bridge config - venet config

Posted by [Vasily Tarasov](#) on Mon, 20 Nov 2006 07:33:58 GMT

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

Hello,

to get rid of venet interface read this thread:

<http://forum.openvz.org/index.php?t=tree&th=270&mid=1516&&rev=&reveal=>

as concerns bridges... This is rather common problem. Read post

[#msg_8380">http://forum.openvz.org/index.php?t=tree&goto=8380&#msg_8380](http://forum.openvz.org/index.php?t=tree&goto=8380) to find out some techniques to solve the problem

HTH,
Vasily.

Subject: Re: Veth on startup (again) + bridge config - venet config

Posted by [sylvain](#) on Mon, 20 Nov 2006 07:56:22 GMT

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

Thanks for your answer.

In both cases this is about modifying the standard script provided by openVZ, that could eventually be broken after an upgrade.

Subject: Re: Veth on startup (again) + bridge config - venet config

Posted by [rickb](#) on Mon, 20 Nov 2006 08:06:44 GMT

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

Welcome to bleeding edge open source!

Subject: Re: Veth on startup (again) + bridge config - venet config

Posted by [duswil](#) on Mon, 20 Nov 2006 08:48:07 GMT

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

Make sure you pay attention to bug #355 as it progresses:

http://bugzilla.openvz.org/show_bug.cgi?id=355

Subject: Re: Veth on startup (again) + bridge config - venet config

Posted by [sylvain](#) on Mon, 20 Nov 2006 09:32:37 GMT

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

Actully if you really look at point one:

Getting rid of venet0

the solution doesn't really work. It prevent the startup script to load and configure venet0.

But there is a mechanism, unknown to me, that copy `"/etc/sysconfig/network-script/ifcfg-venet0"` and `"/etc/sysconfig/network-script/route-venet0"` each time I do `"vzctl start VEID"`

Preventing those script to be copied, and therefore, preventing configuration of venet0 inside the VE would be more helpful than modifying the startup script. Is there a way to do that?

Subject: Re: Veth on startup (again) + bridge config - venet config

Posted by [sylvain](#) on Mon, 20 Nov 2006 10:04:25 GMT

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

Ok I got my answer by looking in the forum:

The variable `OSTEMPLATE=` associated with a proper configuration in `"/etc/vz/dists"` initialize venet0 from inside the VE at startup.

I will change all the script to initialize ethX from inside the VE, creating a new distribution "fake" and setting my OSTEMPLE to this distribution.

When I am done I will provide the script in this forum, hope it could help someone else.

Maybe future release could actually provide for:
distribution-net.conf and distribution-eth.conf that would configure the Ethernet interface instead of the network interface?

I am just making a very modest suggestion with very limited knowledge...

Subject: SOLVED * Veth on startup (again) + bridge config - venet config
Posted by [sylvain](#) on Tue, 21 Nov 2006 03:40:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

This is solved. Here is the step I took to solve this problem. I am aware that bugs have been filed and fix my obsolete my configuration.

This is not too intrusive into openVZ tools and configurations. However it will not appeal to people who wants to control their VE ip configuration from the HN.

First, a bridge start script: /etc/init.d/bridge (to link at proper runlevel)

```
#!/bin/sh  
#
```

```
# Source function library.  
. /etc/init.d/functions
```

```
BRCTL=brctl  
IFCONFIG=ifconfig
```

```
if [ ! -x /usr/sbin/$BRCTL ]; then  
    echo -n $"/usr/sbin/$BRCTL does not exist."; warning; echo  
    exit 0  
fi
```

```
if [ ! -x /sbin/$IFCONFIG ]; then  
    echo -n $"/sbin/$IFCONFIG does not exist."; warning; echo  
    exit 0  
fi
```

```

# Default bridge configuration:
BR_NAME="vzbr0"
BR_ADDRESS="0"
FORWARDING="yes"
PROXY_ARP="yes"

start() {

    echo -n "$Strating bridge $BR_NAME: "

    $BRCTL addbr $BR_NAME
    if [ $? -eq 0 ]; then
    success; echo
    else
    failure; echo; return 1
    fi

    echo -n "$Configuring bridge $BR_NAME: "
    $IFCONFIG $BR_NAME $BR_ADDRESS
    if [ $? -eq 0 ]; then
    success; echo
    else
    failure; echo; return 1
    fi

    if [ $FORWARDING == "yes" ]; then
    echo -n "$Configuring bridge $BR_NAME forwading: "
    if [ -f "$/proc/sys/net/ipv4/conf/$BR_NAME/forwarding" ]; then
        echo 1 > /proc/sys/net/ipv4/conf/$BR_NAME/forwarding
        success; echo
    else
        failure; echo; return 1
    fi
    fi

    if [ $PROXY_ARP == "yes" ]; then
    echo -n "$Configuring bridge $BR_NAME proxy arp: "
    if [ -f "$/proc/sys/net/ipv4/conf/$BR_NAME/proxy_arp" ]; then
        echo 1 > /proc/sys/net/ipv4/conf/$BR_NAME/proxy_arp
        success; echo
    else
        failure; echo; return 1
    fi
    fi

    return 0
}

```

```

stop() {
    echo -n $"Stopping bridge $BR_NAME: "

    $BRCTL delbr $BR_NAME
    if [ $? -eq 0 ]; then
    success; echo
    else
    failure; echo; return 1
    fi

    return 0
}

status() {
    if $BRCTL show 2>/dev/null | grep -q $BR_NAME ; then
    echo $"Bridge $BR_NAME started."
    else
    echo $"Bridge $BR_NAME stopped."
    fi

    return 0
}

restart() {
    stop
    start
}

case "$1" in
    start)
    stop
    start
    RETVAL=$?
    ;;
    stop)
    stop
    RETVAL=$?
    ;;
    restart)
    restart
    RETVAL=$?
    ;;
    status)
    status
    RETVAL=$?

```

```

;;
*)
echo $"Usage: $0 {start|stop|restart|status}"
exit 1
;;
esac

exit $RETVAL

```

I created a new distribution template configuration file: `/etc/vz/dists/nonetdevice.conf`

```

# distribution configuration that inibits:
# * vzctl --addip
# * vzctl --delip
# * vzctl --hostname
#
ADD_IP=null.sh
DEL_IP=null.sh
SET_HOSTNAME=null.sh
SET_DNS=set_dns.sh
SET_USERPASS=set_userpass.sh
SET_UGID_QUOTA=set_ugid_quota.sh
POST_CREATE=postcreate.sh

```

with the corresponding `/etc/vz/dists/scripts/null.sh` script

```

#!/bin/bash

# Null script. It just does not take any actions.

```

Now change the `OSTEMPLATE` variable for you VE in `/etc/vz/conf/X.conf` to be `OSTEMPLATE="nonetdevice"`

I changed a bit `/usr/sbin/vznetcfg`, in the function `init_veth()` only

```

function init_veth()
{
    local dev=$1

    /sbin/ifconfig ${dev} up

    if [ "${VZNET_BRIDGE}" != "" ]
    then
        /usr/sbin/brctl addif ${VZNET_BRIDGE} ${dev}
    fi
}

```


and added the following line in /etc/vz/vznet.conf

```
VZNET_BRIDGE=vzbr0
```

The remaining task consists in configuring the network inside your VE. Enter your VE and, regarding the distribution, setup your configuration files for your Ethernet device (/etc/sysconfig/network-scripts/ifcfg-ethX for centos/redhat/fedora, /etc/network/interfaces for debian based)

This does the job as I wanted.

Subject: Re: Veth on startup (again) + bridge config - venet config

Posted by [duswil](#) on Tue, 21 Nov 2006 04:14:50 GMT

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

Thanks, this looks like it's really useful. Hopefully all will merge into a useful and proper way to do this. Until then, your solution at least appears to be nice. I'll report back to let you know how it works for me as well.

Subject: Re: Veth on startup (again) + bridge config - venet config

Posted by [sylvain](#) on Tue, 21 Nov 2006 06:12:08 GMT

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

Don't hesitate to make amendment to my post. I'd like to hear it's working for you.
