Subject: multiple subnet IP addrss in container
Posted by hkendusers on Wed, 02 Jun 2010 10:39:00 GMT
View Forum Message <> Reply to Message

Hi all,

Since I have two broadband line and wanna split some containers to use broadband A and some containers to use broadband B,
is there anyone know whether possible to assign two different subnet IP addresses to two different containers? Just like the attached photo.

Thx very much!!!

File Attachments
1) vps_diagram.jpg, downloaded 535 times

Subject: Re: multiple subnet IP addrss in container
Posted by kwadrofonik on Fri, 04 Jun 2010 15:18:57 GMT
View Forum Message <> Reply to Message

Openvz will most likely use eth0 for it's venet0 forwarding which is automatic. Merely use --ipadd 10.100.0.1 for the first VE.

For the second VE, route all traffic to 192.168.55.1 to the second VE:

- first enter the second VE and edit /etc/network/interfaces (Debian/Ubuntu). Call eth0 whatever you want the VE's interface name to be. Adjust your netmask accordingly. The gateway should be your router IP address, not the hardware node address.
auto eth0
iface eth0 inet static
address 192.168.55.1
network 192.168.55.0
netmask 255.255.255.0
broadcast 192.168.55.255
gateway 192.168.55.254

- download easymac (www.easyvmx.com) and generate two mac addresses (./easymac.sh -R)
--netif_add (veth name),(generated mac1),(VE's interface name),(generated mac2)
ie --netif_add veth2,00:18:51:4b:f2:fb,eth0,00:18:51:ce:09:36

- create the file /etc/vz/conf/2.mount (where 2 is your second VEID) and make it executable. Enter the following script (credit to original author on this forum). Please research proxy_arp and forwarding which pertains to your distro.

#!/bin/bash
This script source VPS configuration files in the same order as vzctl does

```
# if one of these files does not exist then something is really broken
[ -f /etc/vz/vz.conf ] || exit 1
[ -f $VE_CONFFILE ] || exit 1

# source both files. Note the order, it is important
. /etc/vz/vz.conf
. $VE_CONFFILE

echo waiting for interface
# Configure veth with IP after VPS has started
{
  IP=192.168.55.1
  DEV=veth2
  while sleep 1; do
    /sbin/ifconfig $DEV 0 >/dev/null 2>&1
    if [ $? -eq 0 ]; then
      echo interface found
      /sbin/ip route add $IP dev $DEV
      echo "enabling forwarding"
      /sbin/ifconfig $DEV 0
      echo 1 > /proc/sys/net/ipv4/conf/$DEV/proxy_arp
      echo 1 > /proc/sys/net/ipv4/conf/$DEV/forwarding
      echo 1 > /proc/sys/net/ipv4/conf/eth1/proxy_arp
      echo 1 > /proc/sys/net/ipv4/conf/eth1/forwarding
      break
    fi
  done
} &
```

- if you have two physical interfaces, then outbound routing is also an issue. Using the eth1/proxy_arp (where eth1 is your second interface) should be all you need, but again it may take some playing with.

- Disable any iptable firewall rules while your testing or they'll mess you up.
- When debugging, start with pinging to/from the hardward IP. Then try to/from outside the server.