
Subject: LVS and openvz

Posted by [ebasley](#) on Tue, 16 Dec 2008 15:13:12 GMT

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

got a huge problem with openvz and LVS.

take 2 nodes with Linux version 2.6.26-1-openvz-686 (Debian 2.6.26-11), one (vz1) has ldirectord on CT0 and one VE (ve1), next (vz2) has just one VE (ve2).

if ldirectord (from vz1) is directing traffic to vz2 : all work fine, but if traffic is redirected to local CTO (vz1) kernel thread ksoftirqd takes 100% of the cpu and nothing else happens !

With tcpdump we can see numerous paquets forwarded to ve1, something similar to a loop !!

setup :

VZ1 :

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue state UNKNOWN
    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
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
    link/ether 00:1f:29:0e:f9:4b brd ff:ff:ff:ff:ff:ff
    inet6 fe80::21f:29ff:fe0e:f94b/64 scope link
        valid_lft forever preferred_lft forever
3: br0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UNKNOWN

    link/ether 00:1f:29:0e:f9:4b brd ff:ff:ff:ff:ff:ff
    inet 192.168.2.22/24 brd 192.168.2.255 scope global br0
    inet 192.168.2.206/32 scope global br0
    inet6 fe80::21f:29ff:fe0e:f94b/64 scope link
        valid_lft forever preferred_lft forever
11: veth213.0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UNKNOWN
    link/ether 00:18:51:46:62:00 brd ff:ff:ff:ff:ff:ff
    inet6 fe80::218:51ff:fe46:6200/64 scope link
        valid_lft forever preferred_lft forever
```

/etc/network/interfaces :

auto br0

iface br0 inet static

```
address 192.168.2.22
netmask 255.255.255.0
network 192.168.2.0
broadcast 192.168.2.255
gateway 192.168.2.250
bridge_ports eth0
```

veth, eth0 are bridged in br0.
LVZ VIP (192.168.2.206) is supported by br0

ve1:

```
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
    inet 192.168.2.206/32 scope global lo
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
```

```
3: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UNKNOWN
    link/ether 00:18:51:59:12:98 brd ff:ff:ff:ff:ff:ff
    inet 192.168.2.213/24 brd 192.168.2.255 scope global eth0
    inet6 fe80::218:51ff:fe59:1298/64 scope link
        valid_lft forever preferred_lft forever
```

```
/etc/network/interfaces :
auto lo
iface lo inet loopback
post-up /bin/ip addr add 192.168.2.206 dev lo
```

```
auto eth0
iface eth0 inet static
    address 192.168.2.213
    netmask 255.255.255.0
    gateway 192.168.2.250
dns-nameservers 194.2.0.20 195.68.0.1
dns-search hdf-finance.fr
post-up /sbin/sysctl -w net.ipv4.conf.eth0.arp_ignore="1"
post-up /sbin/sysctl -w net.ipv4.conf.eth0.arp_announce="2"
```

Help will be wellcome.

Thanks

Eric

Subject: Re: LVS and openvz
Posted by [maratrus](#) on Wed, 17 Dec 2008 18:27:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

are you sure that it's OpenVZ specific problem?

Subject: Re: LVS and openvz
Posted by [ebasley](#) on Thu, 18 Dec 2008 16:03:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

may be not.

But the problem may come from the use of veth driver.
In all cases it would be very useful to have an HOWTO about using openvz in LVZ context.

Regards
