

It seems to reason, you should be able to "ping" the bridged device, as it is the gateway out. I use to get "weird" pings, if ProxyARP was not enabled in the kernel on the HN. If that doesn't fix it, I'm assuming, a single interface bridge, and you use VETH devices. I have eliminated a lot of headaches in connecting VE containers to the physical network, by creating static configurations in the VE. In other words, put an "ifcfg-eth0" file, in each VE container, and configure just like, any other physical network adapter. Taking an example from one of my VE containers. I will use 201, as an example. Create a file in a text editor as so.

```
DEVICE=eth0
TYPE=Ethernet
HWADDR=aa:bb:cc:dd:00:11
IPADDR=192.168.245.201
NETMASK=255.255.255.0
ONBOOT=yes
```

It is important to include the auto-gen MAC address for this interface.
Save the file as;

```
/vz/root/201/etc/sysconfig/network-scripts/ifcfg-eth0
```

Restart the VE.

You should now see the IP address I just used.

Now, check the route entries.

Depending on whether you script, to add any routes on the HN for any VE, it should show no default route.

(I do not use any scripts to configure any network device, other than the VETH devices on the HN. Persistent routing and network interfaces can be configured in a proper place, without it)

To add a route to the bridged interface, which should be the gateway out,
In a text editor, open the file;

```
/vz/root/201/etc/rc.d/rc.local
```

If the VE is on the same subnet as the bridge, ADD

```
ip route add default dev eth0
```

Save file.

If you subnet the VE, like I do, ADD;

```
ip route add default eth0
```

```
ip route add 192.168.246.1 via 192.168.254.5
```

This assumes;

```
ip route add <the default route> via <the bridge's IP address>
```

Restart the VE.

Once the VE restarts, you should see the route to the gateway, via the bridge interface.

You should get a successful ping.