
Subject: Ethernet bonding with bridge device on CentOS 5.3

Posted by [batfastad](#) on Tue, 20 Oct 2009 16:50:03 GMT

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

Hi everyone

I'm planning to consolidate a pair of servers onto an existing "beast" box we have and since this beast has 2 GbE sockets and our switch supports dynamic 802.3ad link aggregation with LACP, I'm looking to have those 2 GbE sockets combined into a mode 4 bonded interface.

My first step will be to install the openvz kernel and get that running.

Then before I create any VEs I will get the bond created and working as per this guide...

<http://wiki.openvz.org/Bonding>

Making a slight change to specify mode=4 in modprobe.conf

But how do I then go about combining a veth/bridge with my bond0 interface?

This was the veth tutorial I was planning to follow to setup veths for each VE

http://wiki.openvz.org/Virtual_Ethernet_device#Simple_configuration_with_virtual_Ethernet_device

Then adding each veth into the bridge device

http://wiki.openvz.org/Virtual_Ethernet_device#Virtual_Ethernet_devices_can_be_joined_in_one_bridge

Because these are all "virtual" networking devices, I'm guessing (hoping!) it's something simple like just adding a route for all traffic to go from bond0 to br0 and vice-versa
However I have no idea how to do that

Is any of this correct?

Anyone got any pointers on how to combine my bonded link with a bridge device?

Is it even possible?

Cheers, B

Subject: Re: Ethernet bonding with bridge device on CentOS 5.3

Posted by [ramjet](#) on Tue, 20 Oct 2009 20:03:45 GMT

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

It will work the way you describe. Create a bridge device. Then add your bond0 and the veth devices of your servers to the bridge. Make sure the bridge is up - either `ifconfig br0 0` or `ifconfig br0 up` and enable forwarding for br0 in `/proc`. Never had use for proxyarp. When I tried a coworker with MacOSX complained about errors messages. So be carefull with that.

Btw. are your servers going to run on the same IP subnet?

Are you going to run tagged vlan on top of the bond device? You can. Linux networking is very

flexible.

Have you any iptables rules loaded on the host node? I tend not to as they complicate things.

Feel free to ask more as I'm going to make the same setup Thursday or Friday.

Subject: Re: Ethernet bonding with bridge device on CentOS 5.3

Posted by [batfastad](#) on Tue, 20 Oct 2009 21:29:25 GMT

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

Hi ramjet

Thanks for the info - good news that this is possible!

Yes I want all VEs and the host to be on the same subnet.

No iptables and no tagged VLAN

Yeah I would really appreciate a bit more info on this though if that's ok

Quote:Then add your bond0 and the veth devices of your servers to the bridge. Make sure the bridge is up - either `ifconfig br0 0` or `ifconfig br0 up` and enable forwarding for br0 in `/proc`.

So create the bond per this guide... <http://wiki.openvz.org/Bonding>

Then create the bridge, add VEs to the bridge, add the bond to the bridge

So following this guide... http://wiki.openvz.org/Veth#Virtual_Ethernet_devices_can_be_joined_in_one_bridge

These are the commands I'll need to run:

```
[host-node]# brctl addbr vzbr0
```

```
[host-node]# brctl addif vzbr0 veth102.0
```

```
[host-node]# brctl addif vzbr0 veth103.0
```

```
[host-node]# brctl addif vzbr0 bond0
```

```
[host-node]# ifconfig vzbr0 0
```

```
[host-node]# echo 1 > /proc/sys/net/ipv4/conf/vzbr0/forwarding
```

```
[host-node]# echo 1 > /proc/sys/net/ipv4/conf/vzbr0/proxy_arp
```

But possibly missing the proxy_arp bit... what's the proxy_arp for?

Just read up about it and not sure if I'll need it or not. Doesn't the switch normally deal with arp/mac address/ip translation?

One VE will be a zimbra mail container.

Another will be a Samba/Netatalk/Apache/MySQL container which is our combined NAS/intranet DB system.

Sounds like something like Samba might need proxy_arp enabled for broadcasting whatever it needs to broadcast.

And would I need to do this bit?

http://wiki.openvz.org/Veth#Add_routes_in_CT0

I'm guessing I'd still need to do that but its only for the VE addresses, nothing to do with the bond0

Thanks for all your help... so far

Subject: Re: Ethernet bonding with bridge device on CentOS 5.3

Posted by [maratrus](#) on Sat, 24 Oct 2009 11:25:44 GMT

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

Hello,

Quote:

And would I need to do this bit?

Actually these routes do not affect the network connectivity of a particular CT. So, you may just skip this step.

Subject: Re: Ethernet bonding with bridge device on CentOS 5.3

Posted by [batfastad](#) on Wed, 28 Oct 2009 17:03:20 GMT

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

@maratrus: great thanks for the info!

1) So when I create my bridge I add all my veth<VEID>.x devices to it, and bond0 device.
But do I need to also add eth0 and eth1 to the bridge?
Or by adding bond0 I guess I sort of already add eth0 and eth1 anyway.

2) When adding a veth to a CT using this syntax from the wiki (
http://wiki.openvz.org/Virtual_Ethernet_device#syntax_vzctl_version_.3E_3.0.22)

`vzctl set <CTID> --netif_add <ifname>[,<mac>,<host_ifname>,<host_mac>,<bridge>]`

What do I use for <ifname>?

Should I use bond0?

3) I have an IPCop box acting as our DHCP server.
How do I tell one of the containers to get its IP address from our DHCP server rather than specifying it in the --netif_add syntax above?

Almost there

I'll be trying all this this coming weekend so I'll write up a full report of what I did

Thanks so much for all your help so far

Subject: Re: Ethernet bonding with bridge device on CentOS 5.3

Posted by [maratrus](#) on Fri, 30 Oct 2009 10:45:07 GMT

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

Hello Batfastad,

Quote:

1) So when I create my bridge I add all my veth<VEID>.x devices to it, and bond0 device.

But do I need to also add eth0 and eth1 to the bridge?

Or by adding bond0 I guess I sort of already add eth0 and eth1 anyway.

Yes, you're right. The bond0 device should only be added.

Quote:

What do I use for <ifname>?

Should I use bond0?

You may use whatever you want. It's just for the sake of your convenience. I suppose that eth* would be more convenient for you.

Quote:

How do I tell one of the containers to get its IP address from our DHCP server rather than specifying it in the --netif_add syntax above?

Of course, you don't have to specify IP address in the --netif_add command and obtain it via DHCP server. It may be done in the same way as on a standalone computer.

<http://forum.openvz.org/index.php?t=tree&th=7171&>
