Subject: *SOLVED* OpenVZ and bonding Posted by amorsen on Mon, 04 Jun 2007 12:06:14 GMT View Forum Message <> Reply to Message

I have an existing setup with linux-vserver, where a number of vserver guests each run asterisk. Each guest gets its own VLAN, and routing is then done with policy routing in the host (because vserver does not virtualise routing).

In order to ensure that a failed network interface or switch does not cause an outage, the two ethernet ports have been bonded in the host. I would like to make a similar setup with OpenVZ, without having to do policy routing. It is rather easy without bonding: Create eth0.114, add eth0.114 to guest 114, set up routing and ip addresses in the guest. The same setup with bonding does not work (Adding bond0.114 to the guest fails).

I opened bug report 582 about this; unfortunately it turned out to be a missing feature rather than a bug. Vasily Tarasov wrote:

As remember veterans of the project for some reasons it was deliberately prohibited. In fact, if we move bond device to VE, then all devices, that are slave to this bond should be moved too. On VE stop this "device strucure" should be moved back carefully, etc. So quite a lot of things to think about.

I am not sure why it is necessary to move the slave devices to the guest. In the case where there are multiple VLAN's on top of the bonding device, it is impossible to move the slave devices to the guest -- several guests need them.

Anyway, suggestions are welcome. So far the choice seems to be between giving up on the device failover and giving up on OpenVZ for this particular server.

Subject: Re: OpenVZ and bonding Posted by Andrey Mirkin on Wed, 06 Jun 2007 12:48:57 GMT View Forum Message <> Reply to Message

Well, as I understand you just need a network interface inside VE which is separated from another VE. Right?

Then you just need to create a VE and assign an IP address to venet interface with help of following command:

vzctl set VEID --ipadd IP --save

And you don't need to create vlans at all. Two bonded interfaces will be in host system and you don't need to move them to VE.

Please correct me if I misunderstood your configuration and needs.

Subject: Re: OpenVZ and bonding

Posted by amorsen on Sat, 09 Jun 2007 08:59:08 GMT View Forum Message <> Reply to Message

Unfortunately I need VLAN's. Each VE needs to be on a completely separate network, with different routing. I can do it the linux-vserver way, where I keep all those VLAN's in the host system and just move IP addresses to the VE. Policy routing then sends the traffic in the right directions in the host. This (probably) works, as that is exactly how the server is set up with linux-vserver right now.

It is however very difficult to manage. Only static routing is possible, and all routing changes need to be done in the host, not in the VE. In comparison, we have a router with OpenVZ which provides virtualised routing, complete with BGP. It works very well, but again has the limitation that bonding is not supported.

Subject: Re: OpenVZ and bonding Posted by Andrey Mirkin on Thu, 14 Jun 2007 08:51:29 GMT View Forum Message <> Reply to Message

If you need VLANs in VE then you can do following:

- 1. Create veth device in VE
- 2. Create vlan device in VE based on veth device
- 3. Create bridge in VE0
- 4. Add veth pair device (in VE0) to bridge
- 5. Add eth0 device from VE0 to bridge

Now your VE will have VLAN and other VEs will not be able to see traffic from this VE.

Subject: Re: OpenVZ and bonding Posted by amorsen on Thu, 14 Jun 2007 10:32:18 GMT View Forum Message <> Reply to Message

Thank you for the suggestion Andrey. It is certainly a solution, we might adopt it.

Page 2 of 2 ---- Generated from OpenVZ Forum