
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 structure" 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.
