

---

Subject: Re: dual ethernet woes

Posted by [Steve Hodges](#) on Wed, 15 Aug 2007 03:02:27 GMT

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

---

On 14/08/2007 7:01 PM, Steve Hodges wrote:

> The servers I'm trying to put onto a single hardware node each have 2  
> ethernet interfaces.  
>

OK, after extensive testing I have narrowed down what I believe to be the cause of the problem. The VE always seems to chose the first IP address it has been configured with as the source for new outgoing connections (I have no idea if it does this for existing connections, but I would suspect not).

On the HN if I ping 192.168.205.1, my ping appears to come (by default) from my address on that subnet. I can illustrate that by explicitly setting the source address.

```
ping 192.168.205.1 <-- this works
ping 192.168.205.1 -I 192.168.205.222 <- this works
ping 192.168.205.1 -I 192.168.206.222 <-- this fails
```

```
ping 192.168.206.1 <-- this works
ping 192.168.206.1 -I 192.168.205.222 <- this fails
ping 192.168.206.1 -I 192.168.206.222 <-- this works
```

It seems rather intuitive and reasonable that a machine asked to ping an address that is on the subnet of one of its interfaces will use that interface and set the source address of that ping to the machine's address on the interface it is using.

However, the same does not appear to happen in a VE

```
ping 192.168.205.1 <-- this fails
ping 192.168.205.1 -I 192.168.205.222 <- this works
ping 192.168.205.1 -I 192.168.206.222 <-- this fails
```

```
ping 192.168.206.1 <-- this works
ping 192.168.206.1 -I 192.168.205.222 <- this fails
ping 192.168.206.1 -I 192.168.206.222 <-- this works
```

It looks like the VE is selecting as its source address, the first IP that is set as the source of packets. (In this case the 192.168.206.111 address is set first).

If I swap those around in my script, setting them in the opposite order (by changing the order of the IP addresses in the /etc/vz/conf/111.conf

file) then the behaviour becomes the opposite. That is, the source address now appears to always be 192.168.205.111, and pings on the 192.168.206.0/24 subnet fail.

I guess that the VE has no idea about what subnet it's on. So I guess it can't know that a ping to 192.168.205.1 should more properly come from one or other of its IP addresses. But for correct routing it does make a difference, so how can I force this to happen?

I happen to have a couple of nodes on my network that are far more aware of the network topology (essentially they're my routers). A ping-R from them may be useful to someone who is a better network person than I am. (and I've done this with the IP addresses on the VE set normally).

This is a ping to the gateway on my primary interface (the primary address is set first on the VE)

```
ping -c 1 -R -n 192.168.206.254
PING 192.168.206.254 (192.168.206.254) 56(124) bytes of data.
64 bytes from 192.168.206.254: icmp_seq=1 ttl=63 time=0.333 ms
RR:  192.168.206.111
      192.168.206.222
      192.168.206.254
      192.168.206.254
      192.168.206.222
      192.168.206.111
```

and here is the same thing, but to the gateway on the secondary network

```
ping -c 1 -R -n 192.168.205.254
PING 192.168.205.254 (192.168.205.254) 56(124) bytes of data.
64 bytes from 192.168.205.254: icmp_seq=1 ttl=62 time=0.513 ms
RR:  192.168.206.111
      192.168.205.222
      192.168.205.254
      192.168.205.254
      192.168.206.1
      192.168.206.222
      192.168.206.111
```

192.168.205.254 and 192.168.206.254 are different machines providing a gateway to the internet  
192.168.205.1/192.168.206.1 is the central router that routes between a number of internal networks

If I knew how to use tcpdump I'd probably use that to help diagnose this, but I don't :-(

Steve

---