

Re: Using multiple NICs

Source: <http://linux.derkeiler.com/Newsgroups/comp.os.linux.misc/2006-05/msg00127.html>

- *From:* Robert Heller <heller@xxxxxxxxxxxxx>
 - *Date:* Wed, 03 May 2006 00:10:57 +0200
-

"TH" <noaddress@xxxxxxxxxxxxx>,
In a message on Tue, 2 May 2006 16:04:10 -0500, wrote :

"> Not the answer I was looking for but thanks! :) Had to give it a try.
>
> TH
>
> "Robert Heller" <heller@xxxxxxxxxxxxx> wrote in message
> [news:9ebe3\\$4457c4f6\\$cb248f0\\$15561@xxxxxxxxxxxxxxxxxxxxxxxxxxxxx](mailto:news:9ebe3$4457c4f6$cb248f0$15561@xxxxxxxxxxxxxxxxxxxxxxxxxxxxx)
>> "TH" <noaddress@xxxxxxxxxxxxx>,
>> In a message on Tue, 2 May 2006 13:01:23 -0500, wrote :
>>
>> "> I have this application that transfers huge amounts of data. The data
>> source
>> "> is an NFS share. What I would like to do is mount the NFS share so that
>> the
>> "> reads are from one NIC and move the data out through a different NIC.
>>
>> It is doubtful that you can do **specificly** this.

On re-reading it **sounds** like you might be thinking of having two networks, one with the NFS server and one without the NFS server and a machine with two NICs, one on each network. (I thought at first you wanted to separate NFS reads from NFS writes to/from the same server.) This is trivial to do:

Lets say your NFS server is 192.168.1.100 and is thus on network 192.168.1.0. Your data is meant to land on some machine with an IP of 192.168.2.199 and your compute box has two NICs, eth0 and eth1. eth0 has an IP of 192.168.1.99 and eth1 an IP of 192.168.2.27. You also have two ethernet switches, A and B. You plug the ethernet cable from the NFS server into switch A along with a cable from the compute box's eth0 card. You connect the machine with an IP of 192.168.2.199 to switch B along with a cable from the compute box's eth1 card. You export your NFS share to 192.168.1.99 and mount 192.168.1.100:/whatever on your compute box. The NFS traffic will go through eth0 and on switch A. If your output is headed to 192.168.2.199, it will go through eth1 and through switch B. (It is possible to use a single switch, but I am not sure if you will get burned by network contention — it would depend on

Re: Using multiple NICs

how smart the switch is.)

If you only have one network (eg one network switch) and a machine with two NICs, then what you want to do is bonding. You can have a machine with two IP numbers (on the same network), one per NIC, but things are strange: what would the default route be? With bonding, the kernel merges the traffic on the two NICs into one logical data path, so you have one logical network interface with single IP address and a single route. It just has twice the traffic capacity.

Imagine the difference between a four lane highway, with two lanes in each direction vs two parallel two lane roads (assumes full duplex NICs). The two **separate** roads get you the same place, but you can't shift lanes from one road to the other, should one become congested. Bonding has the advantage of sharing the load on a demand basis. You can't be sure that the NFS read traffic will **exactly** be only 1/2 of the total traffic. It might work out to only be 1/3 of the traffic (or 2/3). With bonding, the difference is made up so that you get to use all of the available bandwidth (assuming your machines can really push that much data). If you do the two network 'trick' above, one NIC could end up being underutilized and the other will be maxed out and total throughput will be less than the max possible. With bonding you get to use ALL of the available bandwidth, even if the traffic is not even.

```
">>
">> ">
">> "> Has anyone tried that before? What do you configure something like
">> that?
">>
">> What you really want to do is called 'bonding', where you bond two (or
">> more) NICs into a single logical NIC. Linux does support this. You
">> need to be sure your EtherNet switch supports it as well. Oh, it really
">> only makes sense if both machines (server AND client) are doing this.
">>
">> You will get the same bandwidth effect as you want, since if you really
">> pull/push data at a rate high enough, the kernel will load share across
">> the two (or more) NICs, although you won't have any one NIC *dedicated*
">> to reads or writes (just like on a SMP system you won't have any one
">> processor dedicated to any specific processing).
">>
">> ">
">> "> Thanks in advance,
">> "> TH
">> ">
">> ">
">> ">
">>
">> Robert Heller -- 978-544-6933
">> Deepwoods Software -- Linux Installation and Administration
">> http://www.deepsoft.com/ -- Web Hosting, with CGI and Database
```

Re: Using multiple NICs

">> heller@xxxxxxxxxxxxx -- Contract Programming: C/C++, Tcl/Tk
">>
">>
">>
">>
">>
">>
">>
">>
">
">
">

Robert Heller -- 978-544-6933
Deepwoods Software -- Linux Installation and Administration
<http://www.deepsoft.com/> -- Web Hosting, with CGI and Database
heller@xxxxxxxxxxxxx -- Contract Programming: C/C++, Tcl/Tk

.