

Re: UDP question.

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-09/msg04435.html>

- *From:* "Holy Aavu" <holyaavu@xxxxxxxx>
 - *Date:* Sat, 16 Sep 2006 13:29:17 +0530
-

On 9/16/06, xixi lii <xixi.limeng@xxxxxxxx> wrote:

2006/9/16, David Schwartz <davids@xxxxxxxxxxxxxxxx>:
>
>> My two adapters has two different IP address, and I bind one IP
>> on one socket,
>> do you mean that I alloc two socket and bind different IP is not
>> helpful?
>
> Correct. You are still sending all the packets *to* the same place.
>
>> In fact, all the packet sent from two socket is go out by one
>> network adapter?
>
> Yes, of course. Why would the kernel send traffic to a destination out an
> interface that doesn't go to that destination?
>
> Suppose you have two interfaces, 1.2.3.4/8 and 10.2.3.4/8, if you are
> sending a packet *to* 1.2.4.5, it will go out the first interface. This
> applies whether the source address is 1.2.3.4 or 10.2.3.4.
>
> By default, the kernel routes traffic based on where it is going, not which
> interface address it came from.
>
> DS
>

Let me explain my network environment, My program is running on a two adapters machine, whose IP is 192.168.0.1/8 and 192.168.0.2/8, then, my destination is two machine, whose IP is 192.168.0.3/8 and 192.168.0.4/8. I use four 100M exchange and a 1000M exchange connected them to ensure the choke point is not at network equipment.

when I use two socket without bonding, one socket is bind 192.168.0.1/8 and sendto 192.168.0.3/8, the other is bind 192.168.0.2/8 and sendto 192.168.0.4/8, but, as you see, I get a result that the speed of send by two adapters is equal to the only one adapter's.

Re: UDP question.

why dont you try a different subnet for the other adapter & its dest.
something like 192.168.1.2/8 & dest 192.168.1.4/8 ?

Because whenever a packet needs to be sent out, the routing table is consulted. (check out the output of route / netstat -nr) So for a given subnet, only the first interface encountered in the table is used. Remaining interfaces are used, only if the first fails fails.

yesterday. I got an uncertain idea, is the problem that IP layer is separate with Eth layer ? when I bind src IP, it just do helpful to IP layer, not real bind the adapter? when I send, the real ethreal adapter is select by IP route? If the two interface can go destination both, IP layer will choose the frist, not use both? Am I right?

yes, that is what the routing table is about.

If so, when I use bonding, the adapter's physical address is the frist one, Do this means that all of the packet come to my machine will go through in the frist one adapter?

No. It just means that the MAC address of the 2 nd card is overridden with the 1st one's MAC. i.e. all packets sent through 2nd NIC will bear the 1st MAC. The reason this is done is not to confuse network equipment which sends out ARP messages etc to resolve IP addresses. (Switches should be intelligent here, else you will have issues with recvng packets on both interfaces).

Thx all.
Best Regard.

xixi

Bhanu.

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx
More majordomo info at <http://vger.kernel.org/majordomo-info.html>
Please read the FAQ at <http://www.tux.org/lkml/>

Re: UDP question.

Re: UDP question.

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>