

## Re: Half Bridge mode and multi-homed box

**Source:** <http://unix.derkeiler.com/Newsgroups/comp.unix.bsd.freebsd.misc/2004-12/0849.html>

---

**From:** jpd (*read\_the\_sig\_at\_do.not.spam.it.invalid*)

**Date:** 12/15/04

Date: 15 Dec 2004 20:12:30 GMT

Begin <32a5p1F3jo17pU1@individual.net>

On 2004-12-15, MattD.. <mattd145@gmail.com> wrote:

- > *On Wednesday 15 Dec 2004 00:23, the whim of a few quarks and leptons to be*
- > *jpd for a while caused the following to be typed:*
- >> *So... it translates PPPoA into "native" ethernet and takes care of*
- >> *"logging in" with the provider?*
- >
- > *Exactly. The router takes all the authentication and link negotiation burden*
- > *from the machine and simply presents it with an Internet routable IP*
- > *address. In theory. :-/*

Well, if it doesn't work as advertised, it might be time to complain to whoever claimed it could do what it doesn't. But first, lessee if there's not some weird problem that can easily be solved. :-)

[snip]

- > *I tried putting the default route that the laptop gets into rc.conf. Same*
- > *result, as was running route flush and adding the route manually. I even*
- > *got to the stage where I set the laptop up, ran netstat -rnfinet,*
- > *disconnected the router and fired the server up and had the output on the*
- > *laptop screen whilst I tried to replicate it on the server.*

Might be, as I mentioned somewhere below, that it does Weird Stuff[tm] with remembering the MAC of whoever told it to open up the connection. Thinking about it again, if it really is (half-a-)bridge-with-featuritis, it might even be the case that the MAC is used to authenticate you on the other side. So you'll really want to flush that (hard reset, call to helldesk of provider) from whoever remembers it for too long. In the most extreme case, you could try and "borrow" the MAC of your laptop on the external interface of your server, but I'd only do that for testing, not "for production".

[snip]

- >> *I'm not sure, but this could very well be how the provider sets it up.*
- >> *I'd probably not do that but I'm not your provider.*
- >
- > *That's not the provider. The provider's gateway is on a completely different*
- > *subnet. I'm still trying to get my head around how the deuce you route an*

> *address on the 82.x.x.x subnet through a gateway on the 204.x.x.x subnet.*

I'd think either something is horribly wrong or you're mixing up stuff. Then again, it just might be a fscked setup. The only reason you really need a gateway IP is so the system knows which MAC to ask for to put in the destination ether address field of packets not for the local subnet (MACs for the local net you can ask the destination directly by arp).

But... I don't know what your edge thingie does: if it has a 204.\* address and a 204.\* gateway on its external side and gives you an 82.\* with an 82.\* gateway for you to use as a gateway, that's just peachy. If it then proceeds and tells you to use a 204.\* for a gateway, that's pretty broken, eh.

> *I*

> *assume it's using PPP tunnels, which is probably where my half bridge is falling down. My machine with the outside IP has no way of getting a packet to the gateway without the modem's assistance via the PPP tunnel because the gateway is on a totally different subnet. The machine also doesn't know the modem can do this, or even that the modem is there. I assume that's why the ADSL router does the "last octet plus one" trick for the gateway. So why does it work on the laptop, this being the case?*

I'd love to see tcpdumps of that, since I haven't a clue what is going on. :-)

[snip!]

> *OK, this is what I've got at the moment. My external address is static and has a PTR to my domain. I've disabled IPv6 on everything, just to be on the safe side (I was running TSPC2 to get a /48 block).*

>

> *The router is doing NAT from the external 82.x.x.x address via 10.0.0.2 to 10.0.0.1. Note NAT, not NAPT, it's port for port, which means I've also got the router set so that 10.0.0.1 is the recipient of everything, a DMZ if you like. There is nothing else, only the router and the server on this subnet connected port to port via a crossover cable.*

Eww. I think I'd prefer a half-bridge mode then. But as I said, if you can put the thing on 192.168.1.1/24 AND have it NAT for 192.168.0.1/23 (note the /23 there) you can kick out the extra natd. But you'll probably want PAT as well, lest you get ``collisions" inbetween the users. Might be that it does that alright but you don't see it for the thing is lazy and only has one client (as far as IT knows).

[snip: repeated application of the same hammer]

> *I was hoping that half-bridge mode on the router was going to make the ADSL connection appear to be just a WAN connection to the server, which indeed it does on the laptop. The differences are these:*

>

> *The server is a gateway. gateway\_enable="YES" in rc.conf.*

> *The server has multiple NICs, on different subnets. The laptop hadn't when I first tried out half bridge.*

Well, you can disable everything while testing, of course. OTOH, I don't think that should be much issue, since all the variable causes is a

```
sysctl net.inet.ip.forwarding=1
```

I momentarily suspected the thing of doing unholy things with caching the mac and Stuff, but that still shouldn't cause your particular problem. You still might to powercycle the thing the hard way before switching from laptop to server in half-bridge mode.

> *The server is running dhcpd and BIND (and a myriad of other services).*  
> *The server is a Dual Xeon SMP box.*  
>  
> *Hmm, look at that last one. I'm wondering if this could be an mpsafenet issue now I type this, since I'm running 5.3-RELEASE and the xe0 adaptor on the laptop is Giant locked and requires debug.mpsafenet=0 in loader.conf.*  
> *Or am I off at tangents? I've just read Robert Watson's post about mpsafenet and various bits of the networking stack being Giant dependent, along with the busdma-SMPng project page, and the more I think about this, the more plausible it seems. I also think it might be a good idea to put the 3Coms back in (see below). The tl driver isn't listed on there as done.*

I wouldn't know, but it's worth a try, I'd think. Maybe someone else in the group can comment or that, or else ask -questions?

[snip: ath0 in hostap mode]

> *I wish I could. It would save me having duplicate lines in /etc/exports for everything, not to mention making my firewall rules a lot simpler. Any advice on this would be VERY helpful if you think it should be possible.*

I misread this at first, I'd assumed we were talking an external AP connected to the server. But you might still be able to pull it off...

> *Last time I tried to bridge the ath0 to the xl0 (internal wired network adaptor) I couldn't ping or communicate with the xl0 (now tl0) address, although I could ping everything else on 192.168.0/24.*

The kernel bridge trick seems to have this problem, whereas building a bridge with netgraph doesn't. You might want to try that, even if it is a bit more involved to setup. I'm not entirely sure it'll work with ethernet<->wifi, but I think it ought to.

[snip: info?]

> *I will do that, as soon as I've slept, and post right back. Another all-nighter, I'm afraid ; -)*

geh.

comp.unix.bsd.freebsd.misc: Re: Half Bridge mode and multi-homed box

> *Thanks for all the advice.*

You're welcome.

--

j p d (at) d s b (dot) t u d e l f t (dot) n l .