

[IPv6] PROBLEM? Network unreachable despite correct route

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-01/msg02645.html>

- *From:* Bernhard Schmidt <berni@xxxxxxxxxxxxxx>
 - *Date:* Tue, 9 Jan 2007 20:36:24 +0100
-

Hi,

I'm having a really ugly problem I'm trying to pinpoint, but failed so far. I'm neither completely convinced it is not related to my local setup(s), nor do I have any clue how this might be caused.

I have several boxes with native IPv6 connectivity at various places. Some of them show symptoms of a lost default route for small periods of time (10–15 seconds several times a day). By symptoms I mean

– traceroute6 from the affected box to any other host dies immediately (the network unreachable does not come from the first hop (the upstream router), but from the local stack itself)

– a local running OpenVPN 2.1_rc1b with UDPv6 transport patched in shows the following output in the syslog file

```
Tue Jan 9 16:48:28 2007 write UDPv6 []: Network is unreachable (code=101)
```

– mtr from the outside to the machine shows that the affected box does not respond anymore, while the hop before (the router) is clean.

– new connects to the box (e.g. ssh) from the outside are stuck (packets get lost, since I'm running my client with tcp_retries=1 I get a timeout)

At the same time, established ssh connections to the box work fine. I can do "ip -6 route" and it shows the default route, both preferred and valid lifetime not exceeded (far from that).

The systems I'm observing this are:

– Dell PowerEdge 750 (P4 with HT), Debian Etch, self compiled kernel 2.6.17.11, connected (e1000) to two upstream Cisco 7200, default route is learned from RIPng (Quagga), static addresses

– Dell OptiPlex GX<something> (P4 with HT, Single Core), SuSE 10.2,

[IPv6] PROBLEM? Network unreachable despite correct route

distribution kernel 2.6.18.5-3-default, connected (tg3) to one upstream Cisco 6500/Sup720, default route learned through stateless autoconfiguration (RA)

– self built AMD Athlon64 (x86_64), Ubuntu Edgy, Distribution kernel 2.6.17-10-generic, connected (forcedeth) to an upstream Linux box (2.6.20-rc3), default route learned through stateless autoconfiguration (RA) as well.

My current believe is that this is an regression introduced in 2.6.17. I have searched for several weeks now why box #1 (the PowerEdge) shows signs of unreachability in the monitoring, but could not find any clue (or verify any reachability problems when I got the monitoring alert). At the same time, a sibling (same hardware, same switch, same network segment, route also learned through Quagga, but different kernel (2.6.16)) of this box did not show any symptoms, so I ruled out the local network.

Also, I upgraded box #2 from SuSE 10.1 (distribution kernel 2.6.16-something) to SuSE 10.2 yesterday. While it was running the OpenVPN/UDPv6 daemon the whole time, there has been exactly one occurrence of the "Network is unreachable" message in the past two weeks before the upgrade (and I can correlate this message with network maintainance where the VPN endpoint was indeed unreachable). Since the upgrade, I have at least 50 lines of that sort in syslog (in about a day).

It is pretty hard to trace this. It seems to appear very seldom, it is not long and I cannot predict the time where it happens by doing more network load or anything else on that machine. IPv4 is fine and without loss in all cases. All network components are dual-stacked, so if there was an L2 issue between the router and the host it would affect IPv4 as well.

Is anyone aware of any issue which might cause this? I've upgraded the PowerEdge to 2.6.19.1 now, but it is too early to tell whether this problem still exists. Does anyone recall a bugreport and maybe a fix for it? A patch or a link to a changeset would be even better, so I could report that to SuSE and Ubuntu to have it included in future kernels.

Thanks,
Bernhard

–

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/>