

Stale IPv6 Neighbor unreachable - Resolve with IPv6 "ARP" equivalent

Right before "mv'ing /dev/myself to /var/home" I came along another odd thing on my border router.

An IPv6 peer was unreachable, i.e. did not respond to ICMP ping, as such the BGP session was down as well. I gave it another indepth look as it happened to be one of our IPv6 upstream peers and had some sort of importance as such.

After talking to the NOC guy of the carrier, which assured that their interface was up, I was a bit confused. He could neither PING my end, nor could I PING his end.

```
#ping ipv6 2001:xxxxxxxx:217
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:xxxxxxxx:217, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

I had however various IPv6 BGP sessions established within the same prefix, so it all seemed ok for the others. For network guys not unusual is however to not only check L3 addressing but also inspect L2 addressing.

IPv4 techs use ARP (Address Resolution Protocol), which maps MAC (Media Access Control) to L3 IPv4 addresses. However this doesn't work, as ARP will only display MAC-IPv4 mappings (addresses obfuscated).

```
#show arp
Protocol Address      Age (min) Hardware Addr  Type  Interface
Internet xxxxxxxx.217    115  0000.0000.00ab  ARPA  GigabitEthernet0/1
Internet xxxxxxxx.2      0  0000.0000.6a9a  ARPA  GigabitEthernet0/3.111
```

So, in the IPv6 world this is a little bit different, as ARP not strictly exists. However, IPv6 uses something similar called "IPv6 neighbor discovery protocol (NDP)". The NDP status can be inspected using the "show ipv6 neighbors" command, which will give you a list of all IPv6 addresses and their MAC address counterparts:

```
#show ipv6 neighbors
IPv6 Address          Age Link-layer Addr State Interface
2001:xxxxxxxxxxxx:4000  0  0009.b766.4000  REACH Gi0/1.10
2001:xxxxxxxxxxxx:9C7D  87  000d.b918.9c7d  STALE Gi0/3.111
2001:xxxxxxxxxxxx:2A80  0  001e.f7f6.2a80  REACH Gi0/3.110
2001:xxxxxxxxxxxx:AB79  0  0013.1937.ab79  STALE Gi0/3.111
```

You may also query a specific entry for a given IPv6 address:

```
#show ipv6 neighbors FE80::xxxxxxxxxxxx:217
IPv6 Address          Age Link-layer Addr State Interface
```

```
2001:xxxxxxxxxxxxxxxx:217          0 0000.0000.00fc STALE Gi0/1
```

As seen in this case, the address in question is in STALE state.

It is beyond the scope of this article to inspect all possible reasons to why this happened.

A good starting place to inspect this is using the "debug ipv6 nd" command, which I may cover in another article.

The most obvious case is as simple as what seldomly happens with IPv4 ARP as well: the L3 and L2 addresses do not match up (maybe due to hardware exchange).

So, in the IPv4-world, you would do something like "clear arp" - but this won't help with IPv6 NDP mismatches.

The simple solution is to clear the IPv6 ND cache:

```
#clear ipv6 neighbors
```

So, if this is really the root of the problem, you may notice a different L2 address afterwards:

```
#show ipv6 neighbors 2001:xxxxxxxxxxxxxxxx:217
IPv6 Address          Age Link-layer Addr State Interface
2001:xxxxxxxxxxxxxxxx:217    0 0000.0000.00ab REACH Gi0/1
```

In my case, the IPv6 peer was reachable by ICMP PING again and as such the BGP session recovered as well.

And now, finally, "mv /dev/myself /var/home/" ;-)