

Extra pitfalls when deploying IPv6 on the Juniper SRX

Preface

After all the trouble I went through to get my Juniper SRX300 working with my Dutch KPN Glasvezel connection in combination with IPv6.. there were still a few things that did not work well...

-I could easily telnet to the ipv6 website of eg google.com on port 80 / 443 and talk to the webserver. So that would imply that all is well, right?.....

→ WRONG.

This is IPv6 so things are still a bit different..

What I noticed was the many IPv6 websites would NOT load.. whereas a telnet to the webserver WOULD work.. it was driving me spare and Wireshark was no help with all those IPv6 addresses I saw... ☹ ☹ ☹

The Internet should either work or not, but not *“for some sites”* right?!

So what to do? Solve problem #1

Of course I googled myself sick. What I at some stage read was that since IPv6 does NOT fragment.. many DNS queries might result in an ICMP **“Packet too large”** message in order for Path MTU to work.

And guess what.. I never saw ICMP “Packet too large” messages on my internal network.. not even if I tested IPv6 with that very useful website <https://test-ipv6.com>.

So.. I reluctantly decided to allow ICMP echo requests and ICMP “packet too large” messages to Enter the WAN interface..

```
match {
  source-address any-ipv6;
  destination-address any-ipv6;
  application [ junos-icmp6-echo-request junos-icmp6-packet-too-big ];
```

That made a LOT of difference.. and more sites were now available.

But... some sites like **www.sixxs.com** would still not load on various browsers and operating systems.. very annoying really..!

Solve problem #2

When closely examining the Router-Advertisement messages within Junos.. I noticed that there is a possibility to make the MTU part of the Router-Advertisement messages. (RA) And guess what... since IPv6 does NOT fragment... it is VERY picky when it comes to Path-MTU discovery. And test-ipv6.com was still complaining on a regular bases (depending on the country you picked to perform the test with..) that there was SOMETHING wrong with my DNS resolving of large packages.

Of course I tried many resolvers. (google, OpenDNS, my ISP, IPv4 / IPv6, you name it)

Extra pitfalls when deploying IPv6 on the Juniper SRX

Then I decided..... to *make the MTU part of the RA message*:

So now my router-advertisement protocol configuration looked like this:

```
interface irb.100 {
  managed-configuration;
  other-stateful-configuration;
  link-mtu;
  prefix 2a02:xxxx:xxxx:xxx::/64 {
    on-link;
    autonomous;
  }
}
```

However..... scrupulous Wireshark research revealed that the RA was advertising the MTU to be 1500 Bytes, whereas my WAN connection (fiber) was using PPPoE so here the MTU was 1492B.

So..... I then did something that looked for me to be very counter intuitive:

Since you can NOT set the MTU size on the router-advertisement, I changed the MTU on the *VLAN interface IRB.100* to be 1492 B, just like my WAN interface.

```
SRX300# top show interfaces irb.100 family inet6
mtu 1492;
```

Wireshark showed that the proper MTU was not being advertised towards the clients.

And lo and behold.... The sites that would previously NOT load were loading!!!!

Does that now finally solve Everything and is IPv6 superior?..

I wish it was.. alas. There are some sites,- the Dutch news site **nopstart** for instance -, that via geofencing wants to make sure that the client IPv6 address is truly local. And for some reason.. they just don't like my (Dutch..) IPv6 prefix. So for that particular site.. I have created a firewall rule that simply blocks IPv6 access. Sigh..

Quite the Adventure..

I wish I could say that IPv6 on the Juniper SRX was easy.. but it was not. The documentation is very insufficient. It is all trial and error. And the ISPs don't help as they do not publish the requirements for IPv6. (Prefix Delegation, how to find your next-hop, how to find your prefix, requirements for DNS resolving, ICMP settings etc etc.)

The most useful Junos command I found that gave me most information was;

```
SRX300# run show dhcpv6 client binding detail
```

```
Client Interface/Id: pp0.0
  Hardware Address:      d8:xx:xx:xx:xx:xx
  State:                 BOUND(DHCPV6_CLIENT_STATE_BOUND)
  ClientType:           STATEFUL
  Lease Expires:        2022-09-09 21:33:30 CEST
  Lease Expires in:     252098 seconds
```

Extra pitfalls when deploying IPv6 on the Juniper SRX

```
Lease Start:                2022-09-06 21:33:30 CEST
Bind Type:                  IA_PD
Preferred prefix length    0
Sub prefix length          0
Client DUID:                LL0xxx-xx:xx:xx:xx:xx:xx
Rapid Commit:              Off
Server Identifier:         fe80::xxxx:xxx:xxx:xxxx
Update Server              Yes
Client IP Prefix:          2a02:xxxx:xxxx::/48
```

DHCP options:

```
Name: server-identifier, Value: LL_TIME0x6-0x6c79e-xx:xx:xx:xx:xx:xx
Name: dns-recursive-server, Value: 2a02:xxxx:xxxx:xx,xxxx:xxxx:xxxx:xx
```

Update RA interfaces:

```
Interface: irb.200
RA Prefix:      2a02:xxxx:xxxx::/64
```

This command showed me my PD prefix, as well as the next hop.

It has been a long journey.. and IPv6 has proven to be a fierce adversary.

Now I am going to watch some Netflix.. On IPv6!!

If this document has helped you with your struggles.. send me a message, I'd like to hear from you.