Preface

The documentation on this subject, I find, is sketchy to say the least. Many options are not well explained or documented and as IPv6-is-IPv6.. everything seems just a bit harder. Since I struggled to come to terms with this subject in the past myself, now's the time to document my findings.

Topology

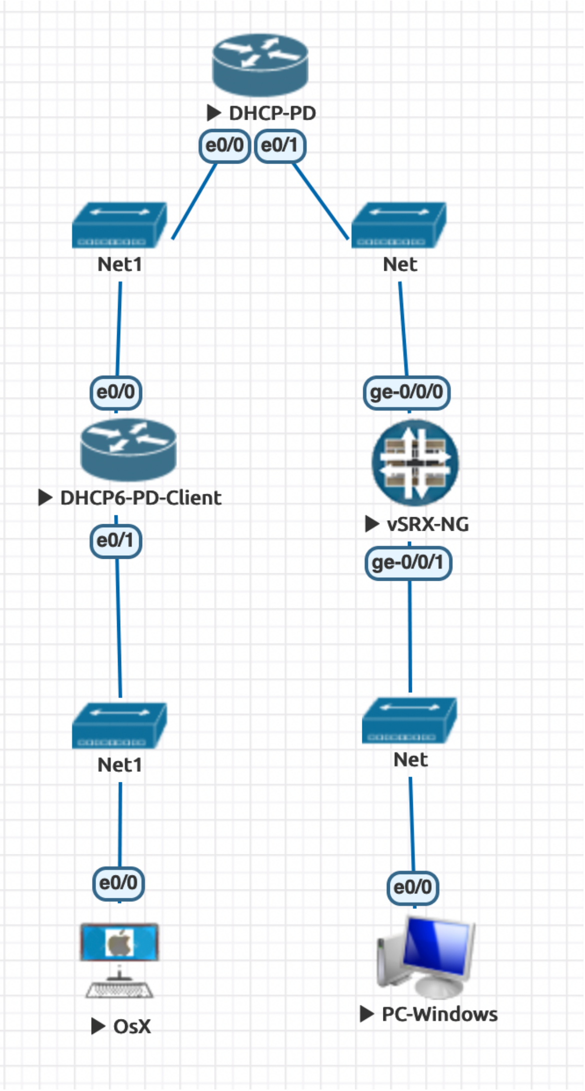

DHCP-PD is the IPv6 DHCP server that serves all /48 Prefixes as opposed to individual networks. The individual /64 subset-networks are allocated and announced by the clients **DHCP6-PD** and **vSRX-NG** via SLAAC towards the "PC clients" OsX and PC-Windows.

The "home routers/firewalls" DHCP6-PD-client (IOS) and vSRX-NG (Junos) will receive a prefix on their respective Internet facing interfaces `e0/0` and `ge-0/0/0`.

The remarkable thing is that…. **BOTH DEVICES will NOT configure a public IPv6 address on their Internet facing Interfaces**! Just the link-local FE80::/10 address will do..

Which of course works fine, but can just not be used as an IPv6 source address for routing.

Configuration of the DHCP-PD server

Enable IPv6 unicast routing and create a DHCPv6 pool to delegate the Prefix with:

```
      ipv6 unicast-routing
      ipv6 dhcp pool IPV6-BASE
            prefix-delegation pool DHCPV6-POOL1 lifetime 1800 600
```

The local pool that Prefixes will be delegated from is:

```
      ipv6 local pool DHCPV6-POOL1 2001:DB8:1200::/40 48
```

The configuration of the interface towards the Cisco client is:

```
      interface Ethernet0/0
         description Link to Cisco DHCP6-PD-Client
         no ip address
         ipv6 address 2001:A:B:C:4:3:2:1/64
         ipv6 enable
         ipv6 dhcp server IPV6-BASE
```

So here /48 Prefixes can be handed out, from the range of **2001:DB8:1200::/40**.

On the Server, (once the client is configured) the handing out of the prefix towards the Cisco client via DHCPv6 can be observed:

```
DHCP-PD# show ipv6 dhcp binding
[…]
Client: FE80::A8BB:CCFF:FE00:2000
  DUID: 00030001AABBCC002000
  Username : unassigned
  VRF : default
  Interface : Ethernet0/0
  IA PD: IA ID 0x00030001, T1 300, T2 480
    Prefix: 2001:DB8:1201::/48
            preferred lifetime 600, valid lifetime 1800
            expires at Sep 15 2022 11:42 AM (1672 seconds)
```

Observer the Client Link-Local address, DUID[1], outgoing Interface and the Prefix. These are all part of the Cisco client.

Configuration of the Cisco DHCP6-PD-Client

The Cisco client configuration looks as follows: only the link-local address is used on the Internet facing interface and the “dhcp pd” client is used to receive the delegated Prefix from the server. The received Prefix is associated with the name “**ISP-PREFIX**”:

```
      ipv6 unicast-routing
      interface Ethernet0/0
         no ip address
         ipv6 enable
         ipv6 dhcp client pd ISP-PREFIX
```

What does the Interface look like after configuration?

[1] DHCP Unique Identifier ID

```
DHCP6-PD-Client(config)# do sh ipv6 int br
Ethernet0/0            [up/up]
    FE80::A8BB:CCFF:FE00:2000
```

Since SLAAC (=Autoconfiguration) is NOT used on this interface, we need to supply the device with a fully qualified, static default route:

```
ipv6 route ::/0 Ethernet0/0 FE80::A8BB:CCFF:FE00:1000
```

Here the next-hop is of course the link-local address of the server. In the case of a **PPP**(oe) interface (as I am using in my network), **the next-hop link local address is not required**.

We know the next-hop IPv6 from the PD announcement:

```
DHCP6-PD-Client# show ipv6 dhcp interface e0/0
Ethernet0/0 is in client mode
  Prefix State is OPEN
  Renew will be sent in 00:01:53
  Address State is IDLE
  List of known servers:
    Reachable via address: FE80::A8BB:CCFF:FE00:1000
    DUID: 00030001AABBCC001000
    Preference: 0
    Configuration parameters:
      IA PD: IA ID 0x00030001, T1 300, T2 480
        Prefix: 2001:DB8:1201::/48
                preferred lifetime 600, valid lifetime 1800
                expires at Sep 15 2022 11:52 AM (1614 seconds)
      Information refresh time: 0
  Prefix name: ISP-PREFIX
  Prefix Rapid-Commit: disabled
  Address Rapid-Commit: disabled
```

Now we can hand-out prefixes within the /48 range on downstream interfaces:

```
interface Ethernet0/1
 no ip address
 ipv6 address ISP-PREFIX ::1:0:0:0:1/64
 ipv6 enable
```

And here comes the Magic: `ISP-PREFIX` `2001:DB8:1201::/48` is appended with `::1:0:0:0:1/64` for the first interface / clients. So ALL clients that now use SLAAC and are connected to `Ethernet0/1` will create IPv6 addresses within the range
`2001:DB8:1201:1::/64`

Mind you: as opposed to Junos, SLAAC is automatically enabled on Cisco devices once a working IPv6 address is configured on a Router.

The Client PC

The effect on an “OsX PC” (=just another Cisco IOS device really) is now:

```
Ethernet0/1            [up/up]
    FE80::A8BB:CCFF:FE00:2010
    2001:DB8:1201:1::1
```

And since the “PC” uses SLAAC, it has also received the Default Gateway / Router:

```
DHCP6-PD-Client# sh ipv6 route
[…]
```

```
S   ::/0 [1/0]
     via FE80::A8BB:CCFF:FE00:1000, Ethernet0/0
```

Configuration of the Juniper vSRX

This configuration is a bit more complicated. Not just because it's Junos,- which always requires a bit more work-, but simply as well because we're dealing with a Firewall here. The extra security layer of a Firewall does tend to complicate matters. (ICMP for Neighbor / Router discovery etc..)

The configuration on the Internet facing Interface of the Firewall is:

```
root@vSRX-1# show interfaces ge-0/0/0
unit 0 {
    family inet {
        address 10.0.0.2/24;
    }
    family inet6 {
        dhcpv6-client {
            client-type stateful;
            client-ia-type ia-pd;
            client-identifier duid-type duid-ll;
            req-option dns-server;
            retransmission-attempt 9;
            update-server;
```

The interface is a STATEFUL (NOT Autoconfig!) client, and the client-ia-type (= identity association type) is Identity association for prefix delegation.
Since that is what we want to achieve.

The DHCP Unique Identifier (DUID) is the Link-Layer address = MAC address. The statement "**update-server**" refers to the fact that the Prefix learned on this Interface, itself is forwarded to the DHCP server that will be used for SLAAC on this Firewall. As of course (just like with the Cisco), the device is a Client of the Server, and Serves the clients the /64 prefix.

How will the Juniper perform SLAAC? For that it uses "**router-advertisements**":

```
root@vSRX-1# show protocols
router-advertisement {
    interface ge-0/0/1.0 {
        managed-configuration;
        other-stateful-configuration;
        link-mtu;
        prefix 2001:db8:1200:3::/64;
```

So here the prefix `2001:db8:1200:3::/64` will be announced via SLAAC to the clients. The Security settings of the SRX will be kept outside of the discussion; just make sure you have policies to traverse from your zone to the untrust zone, and allow ICMPv6 on your internet facing interface.

Of course as before, we need a static default route on the Juniper:

```
root@vSRX-1# show routing-options
rib inet6.0 {
    static {
        route ::/0 {
            qualified-next-hop fe80::a8bb:ccff:fe00:1010 {
                interface ge-0/0/0.0;
```

How can we find out the next-hop IPv6 address of the “ISP”? This is received from the PD messages of the Server:

```
root@vSRX-1# run show dhcpv6 client binding detail

Client Interface/Id: ge-0/0/0.0
     Hardware Address:             50:00:00:04:00:01
     State:                        BOUND(DHCPV6_CLIENT_STATE_BOUND)
     ClientType:                   STATEFUL
     Lease Expires:                2022-09-15 11:31:53 UTC
     Lease Expires in:             1749 seconds
     Lease Start:                  2022-09-15 11:01:53 UTC
     Bind Type:                    IA_PD
     Preferred prefix length       0
     Sub prefix length             0
     Client DUID:                  LL0x3-50:00:00:04:0a:01
     Rapid Commit:                 Off
     Server Identifier:            fe80::a8bb:ccff:fe00:1010
     Update Server                 Yes
     Client IP Prefix:             2001:db8:1200::/48

DHCP options:
    Name: server-identifier, Value: LL0x1-aa:bb:cc:00:10:00
```

The /48 Prefix is announced by the SRX:

```
root@vSRX-1# run show ipv6 router-advertisement prefix
2001:db8:1200:3::/64
Interface: ge-0/0/1.0
  Advertisements sent: 24, last sent 00:05:05 ago
  Solicits sent: 1, last sent 02:17:38 ago
  Solicits received: 0
  Advertisements received: 48
  Solicited router advertisement unicast: Disable
  IPv6 RA Preference: DEFAULT/MEDIUM
  Advertisement from fe80::a8bb:ccff:fe00:5000, heard 00:01:24 ago
    Managed: 0 [1]
    Other configuration: 0 [1]
    Link MTU: 1500 bytes
    Reachable time: 0 ms
    Default lifetime: 1800 sec
    Retransmit timer: 0 ms
    Current hop limit: 64
```

The Client PC

The effect on the client:

```
PC-Windows# show ipv6 interface brief
Ethernet0/0            [up/up]
    FE80::A8BB:CCFF:FE00:5000
    2001:DB8:1200:3:A8BB:CCFF:FE00:5000

PC-Windows# show ipv6 route
[…]
ND  ::/0 [2/0]
     via FE80::5200:FF:FE04:2, Ethernet0/0
```

The proof is in the pudding

If all goes well.. (and all security policies are in place..) we should now be able to ping between the 2 PC's across this IPv6 "internet":

```
PC-Windows# ping 2001:DB8:1201:1::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:1201:1::1, timeout is 2
seconds:
!!!!!
```

And so.. IPv6 works once more.. **with Prefix Delegation**!

Pittfalls in IPv6

There are multiple ones.. ready my other Document on Ipv6 and DNS Queries in combination with ICMP "packet too large" messages.. as IPv6 will not allow fragementation, it uses Path-MTU and relies HEAVILY on ICMP in combination with "packet too large" messages.. Oh joy..