A study in Neighbor Discovery (ND) and IPV6

-ND is the replacement for ARP in IPv6 by using ICMP types 135 for Neighbor Solicitation and 136 for Neighbor Discovery.
-ND is used for Router Solicitations and Router Discovery by means of ICMP types 133, 134, 137.

Note: The example of a Router Solicitation below, have been captured with:

```
# tcpdump -i en0  '(icmp6 && (ip6[40] >= 133 && ip6[40] <= 137))'
```

In the examples below, the following Router – and Desktop Ipv6 Link Local addresses are being used:

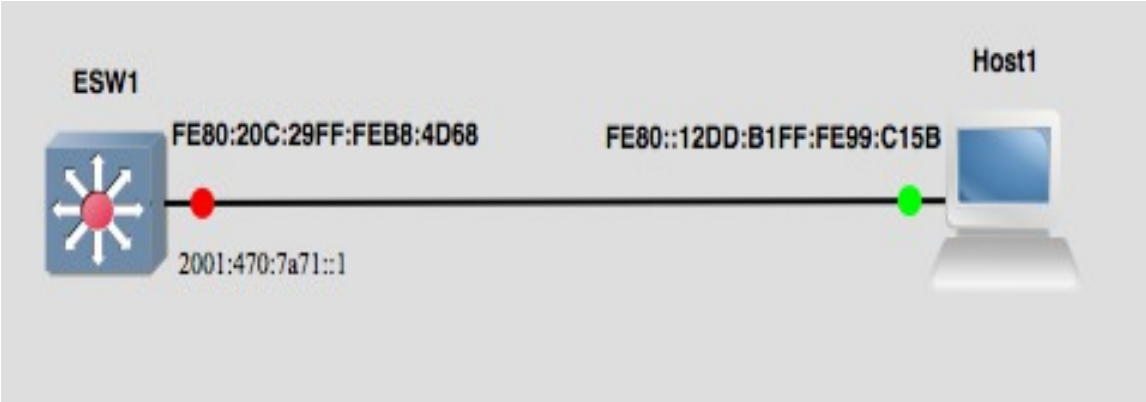

Step 1: Router Solicitation, or querying for the IPv6 Prefix

The client (FE80::12DD:B1FF:FE99:C15B) solicits for a Routing Prefix by **ICMP type 133**, using its FE80 LinkLocal address as source and FF02::2 as the Multicast Destination:

```
▶ Frame 6: 62 bytes on wire (496 bits), 62 bytes captured (496 bits)
▶ Ethernet II, Src: Apple_99:c1:5b (10:dd:b1:99:c1:5b), Dst: IPv6mcast_02 (33:33:00:00:00:02)
▼ Internet Protocol Version 6, Src: fe80::12dd:b1ff:fe99:c15b, Dst: ff02::2
    0110 .... = Version: 6
  ▶ .... 0000 0000 .... .... .... .... .... = Traffic class: 0x00 (DSCP: CS0, ECN: Not-ECT)
    .... .... .... 0001 0110 0101 1100 0100 = Flowlabel: 0x000165c4
    Payload length: 8
    Next header: ICMPv6 (58)
    Hop limit: 255
    Source: fe80::12dd:b1ff:fe99:c15b
    [Source SA MAC: Apple_99:c1:5b (10:dd:b1:99:c1:5b)]
    Destination: ff02::2
    [Source GeoIP: Unknown]
    [Destination GeoIP: Unknown]
▼ Internet Control Message Protocol v6
    Type: Router Solicitation (133)
    Code: 0
    Checksum: 0xf864 [correct]
```

Step 2: The Router Advertisement, or Obtaining the iPv6 prefix

The Router Advertisement / response comes in the form on a **ICMP type 134** Datagram to Multicast Group FF02::2. The Source is the Router with Link Local address FE80:20C:29FF:FEB8:4D68 :

```
▼ Ethernet II, Src: Vmware_b8:4d:68 (00:0c:29:b8:4d:68), Dst: IPv6mcast_01 (33:33:00:00:00:01)
  ▼ Destination: IPv6mcast_01 (33:33:00:00:00:01)
      Address: IPv6mcast_01 (33:33:00:00:00:01)
      .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory default)
      .... ...1 .... .... .... .... = IG bit: Group address (multicast/broadcast)
  ▼ Source: Vmware_b8:4d:68 (00:0c:29:b8:4d:68)
      Address: Vmware_b8:4d:68 (00:0c:29:b8:4d:68)
      .... ..0. .... .... .... .... = LG bit: Globally unique address (factory default)
      .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: IPv6 (0x86dd)
▼ Internet Protocol Version 6, Src: fe80::20c:29ff:feb8:4d68, Dst: ff02::1
    0110 .... = Version: 6
  ▼ .... 0000 0000 .... .... .... .... .... = Traffic class: 0x00 (DSCP: CS0, ECN: Not-ECT)
      .... 0000 00.. .... .... .... .... .... = Differentiated Services Codepoint: Default (0)
      .... .... ..00 .... .... .... .... .... = Explicit Congestion Notification: Not ECN-Capable Transport (0)
    .... .... .... 0000 0000 0000 0000 0000 = Flowlabel: 0x00000000
    Payload length: 136
    Next header: ICMPv6 (58)
    Hop limit: 255
    Source: fe80::20c:29ff:feb8:4d68
    [Source SA MAC: Vmware_b8:4d:68 (00:0c:29:b8:4d:68)]
    Destination: ff02::1
```

As part of the Router Advertisement in **ICMP Type 134**, the IPv6 PREFIX is published:

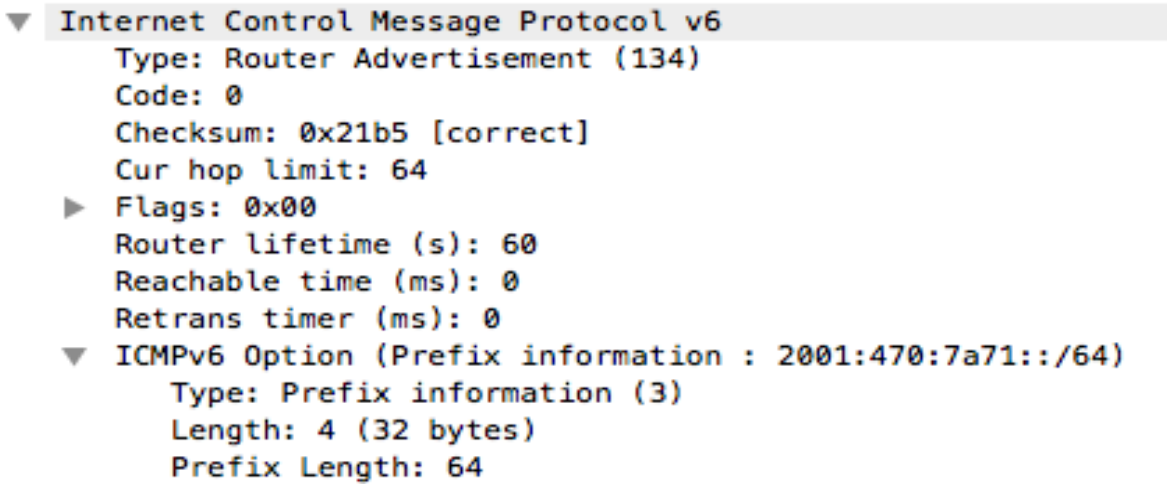

So the /64 Prefix published is: 2001:470:7a71::/64. The client will use EUI-64 in this case to configure its entire IP address which amounts to: **`2001:0470:7a71:0000:12dd:b1ff:fe99:c15b`**. Here **`2001:0470:7a71:0000`** `is the /64 bit Network prefix and` **`12dd:b1ff:fe99:c15b`** `i`s the 64 bit Host – ID.

The router ALSO declares itself as DEFAULT router (::) :

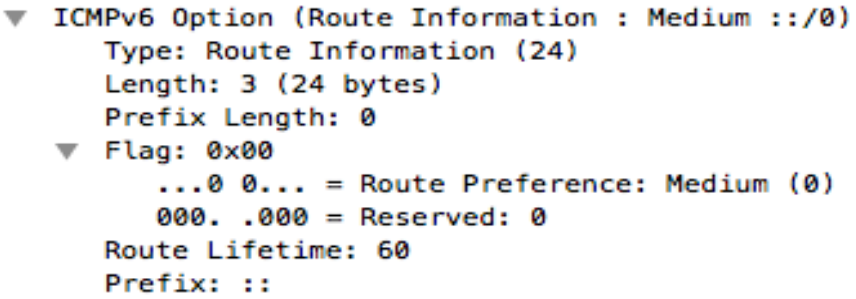

As part of the Advertisement, the DNS server(s) are Published:

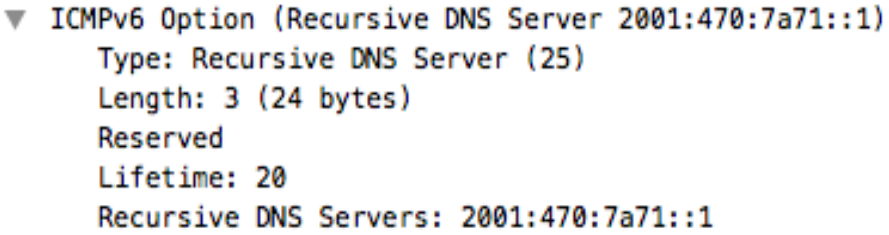

As well as the DNS search list:

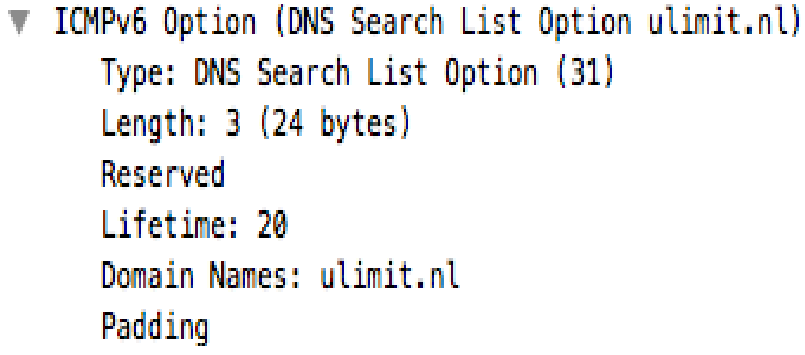

Step 3: Neighbor[1] Discovery / Solicitation

Here the Client 2001:0470:7a71:0:12dd:b1ff:fe99:c15b will try to discover the MAC address of the Router at IPv6 Global Unicast Address 2001:470:7a71::1 to Multicast Address FF02::1:FF00:1 :

```
▼ Internet Protocol Version 6, Src: 2001:470:7a71:0:12dd:b1ff:fe99:c15b, Dst: ff02::1:ff00:1
    0110 .... = Version: 6
  ▶ .... 0000 0000 .... .... .... .... .... = Traffic class: 0x00 (DSCP: CS0, ECN: Not-ECT)
    .... .... .... 0000 0000 0000 0000 0000 = Flowlabel: 0x00000000
    Payload length: 32
    Next header: ICMPv6 (58)
    Hop limit: 255
    Source: 2001:470:7a71:0:12dd:b1ff:fe99:c15b
    [Source SA MAC: Apple_99:c1:5b (10:dd:b1:99:c1:5b)]
    Destination: ff02::1:ff00:1
    [Source GeoIP: Unknown]
    [Destination GeoIP: Unknown]
```

The IPv6 address that is to be resolved (**2001:470:7a71::1**), is found in the ICMP Type 135 datagram:

```
▼ Internet Control Message Protocol v6
    Type: Neighbor Solicitation (135)
    Code: 0
    Checksum: 0x3333 [correct]
    Reserved: 00000000
    Target Address: 2001:470:7a71::1
```

Step 4: The Neighbor Advertisment answer (simular to the ARP response)

Here is the Router's response:

```
Source: 2001:470:7a71::1
Destination: 2001:470:7a71:0:12dd:b1ff:fe99:c15b
[Destination SA MAC: Apple_99:c1:5b (10:dd:b1:99:c1:5b)]
```

The MAC address is found in the **ICMP Type 136** Datagram:

1 Sorry for the American spelling, I didn't come up with this

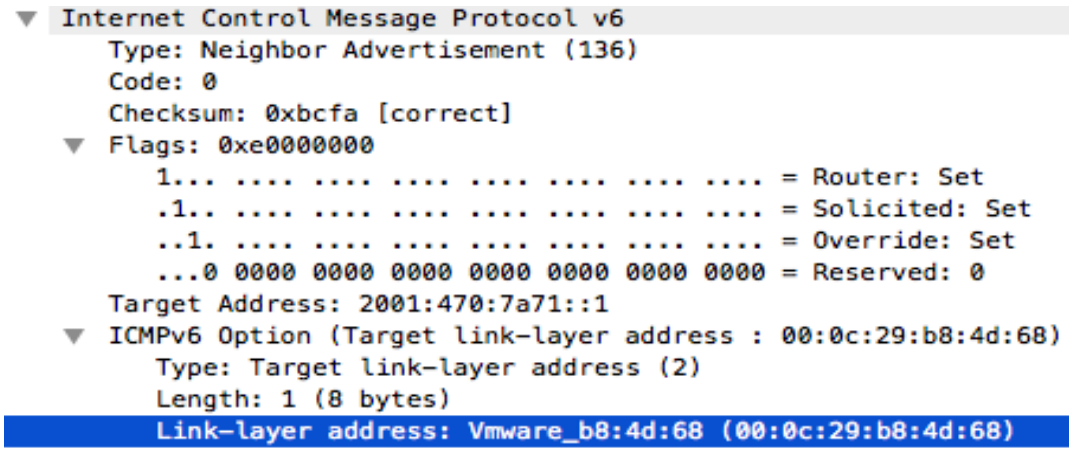

The client now has obtained its usefull IPv6 Global Unicast Address as well as the default (::) Router. I hope this DE-mistifies IPv6 Neighbor Discovery a bit..

---->>>> Let the Games begin.. <<<<----