pfSense - Bug #4206

Missing route creation on DHCP-PD lease where ia-na != ia-pd

01/12/2015 03:52 PM - Anders Lind

Status:	Resolved	**Start date:**	01/12/2015
Priority:	Low	**Due date:**	
Assignee:	Chris Buechler	**% Done:**	100%
Category:	DHCP (IPv6)	**Estimated time:**	0.00 hour
Target version:	2.3		
Affected Version:	All	**Affected Architecture:**	

Description

The long story (https://forum.pfsense.org/index.php?topic=86374.0) short:
When a prefix is delegated through DHCP for IPv6 from a pfSense edge router to a SOHO D-Link sub-router then only the DHCPv6 info is supplied to that sub-router, but pfSense does not create the route to the sub-router/delegated sub-network.

The problem short:
The route is never created because the resulting string generated by /usr/local/sbin/prefixes.php is:
/sbin/route change -inet6 <delegated network> <but missing WAN address of the sub-router>
E.g.:
/sbin/route change -inet6 2a02:abcd:dcba:3fff::/64

The cause (https://forum.pfsense.org/index.php?topic=86374.msg474928#msg474928):
The lease file /var/dhcpd/var/db/dhcpd6.leases contains for the sub-router different strings for ia-na and ia-pd that leads to that the WAN address of the sub-router becomes the empty string, because /usr/local/sbin/prefixes.php at lines:

```
     55 $routes = array();
     56 foreach ($duid_arr as $entry) {
     57         if(!empty($entry['ia-pd'])) {
     58                 $routes[$entry['ia-na']] = $entry['ia-pd'];
     59         }
     60 }
```

, fails, because the IPv6 address of ia-na and the IPv6 network of ia-pd lies in each (but different) entries of the $duid_arr (and not the same entry if the strings of ia-na and ia-pd had been equal - see below!)

This happens because:

```
ia-na:
ia-na "\273\240\300\034\000\003\000\001\300\240\273\034\xxx\xxx" {
ia-na in hex:
         BB   A0   C0   1C   00   03   00   01   C0   A0   BB   1C   XX   XX

ia-pd:
ia-pd "\000\000\000\000\000\003\000\001\300\240\273\034\xxx\xxx" {
ia-pd in hex:
         00   00   00   00   00   03   00   01   C0   A0   BB   1C   XX   XX

mac address of the sub-router:           c0 :a0 :bb :1c :xx :xx
```

and because the start of the prefixes.php file:

```
     10 $duid_arr = array();
     11 while (( $line = fgets($fd, 4096)) !== false) {
     12         // echo "$line";
     13         if(preg_match("/^(ia-[np][ad])[ ]+\"(.*?)\"/i", $line, $duidmatch)) {
```

```
    14                  $type = $duidmatch[1];
    15                  $duid = $duidmatch[2];
    16                  continue;
    17          }
```

, stores the entire octet strings into $duid instead of e.g. only the "real" duid part meaning excluding the first 4 octets. Well the DUID in this example is a type 3 (DUID-LL).

Question 1 is why this difference? I do not have a clue e.g. if it is the dhcp6 server or the client (the D-Link sub-router) that makes the alternative ia-na string compared to the ia-pd string, but if we look at the MAC address:

```
c0 :a0 :bb :1c :xx :xx <-- MAC
c0 :a0 :bb :1c         <-- First 4 blocks of MAC
BB  A0  C0  1C         <-- two blocks (first and third) switch places
BB  A0  C0  1C == \273\240\300\034\ <-- start/first 4 octets of the ia-na string.
```

Question 2: Are the possible solutions to fix prefixes.php (line 13 and line 15) by either:
1) taking a part of the ia-na string and ia-pd string that corresponds to the DUID and remove(/leave out) the first 4 blocks (1 block => \xxx) or
2) forcing the first 4 blocks zeroed out (\000) or
3) do and verify what the dhcpv6 service or my D-Link sub-router does?

I just guess it has nothing to do with RFC6355 (http://tools.ietf.org/html/rfc6355), but "just" RFC3315 (http://tools.ietf.org/html/rfc3315) combined with some ISC DHCPv6 stuff and/or the D-Link stuff of which I have no understanding.

Associated revisions

Revision 2caea6a2 - 01/18/2016 01:50 PM - Anders Lind

Bugfixes & handling $duid and $type, Fixes #4206

This patch addresses:
1. Handling of IA_NA and IA_PD strings (that contain IAID+DUID content) using only the DUID part.
2. Fixing regular expression matching with respect to the IAID+DUID string regarding the legal \" substring (used in ISC DHCPv6 leases).
3. Checking the $duid variable before use. Default case for $type in the switch case statement.

Please see the ticket for further information.

History

#1 - 01/13/2015 12:10 AM - Jim Thompson

- Assignee set to Chris Buechler

#2 - 01/13/2015 12:36 AM - Chris Buechler

- Affected Version changed from 2.2 to All

#3 - 01/13/2015 01:50 AM - Chris Buechler

- Subject changed from Missing route creation on DHCP-PD lease to Missing route creation on DHCP-PD lease where ia-na != ia-pd

- Status changed from New to Confirmed

- Target version changed from 2.2 to 2.2.1

updated subject to root cause of issue.

Anders: asked about getting a pcap of the DHCPv6 traffic in your forum thread, that'd be helpful to verify what's on the wire vs. what ends up in dhcp6.leases.

Clients can set their IA_PD and can have multiple ones configured, such as one per interface. That's probably why you're ending up with part of the MAC there.

This is an unusual edge case in a relatively little-used feature and something that's never worked. Going to have to push to 2.2.1 at this late stage in 2.2. PD routes in general work fine, it's just the circumstance in subject that's an issue.

#4 - 01/20/2015 02:21 PM - Anders Lind

- File packetcapture.cap added

- File Status DHCPv6 leases.png added

- File dhcpd6.leases added

Here is a follow-up with 3 attachments: The pcap file, a screenshot of the "Status: DHCPv6 leases" page and the dhcpd6.leases file.

Before I started the capture I changed the MAC addresses of my laptop's LAN interface (attached to pfSense LAN), the LAN interface of pfSense and the WAN interface of the D-Link sub-router (attached to pfSense LAN.)
It is a UDP capture on the pfSense LAN interface for IPv6 packets only made in promiscuous mode.
MACs:
Laptop: 00:11:22:00:11:22
pfSense LAN: 00:44:33:44:33:00
Sub-router/D-Link WAN: 00:11:22:33:44:55

The capture shows what happens when the D-Link router is connected to the pfSense LAN. (Starting with packet number 29.)

Just some comments and observations:
1. I guess that D-Link uses the same code base for their SOHO routers, so I guess it is likely that handling ia-na != ia-pd is also an issue on other of their router models (and maybe brands of other SOHO routers as well.)

2. I am unable to alter e.g. the DUID in the D-Link router. I guess it is because that the router like other SOHO routers are targeted the mainstream home users and has to be easy to use, so D-Link cuts the "superfluous" stuff away. Therefore till now there is very little people can do besides setting up a static IPv6 environment.

3. I wonder why the DUID described in the pfSense GUI tells the DUID is "33:00:03:00:01:00:11:22:22:33:44:55". *)
I guess it is the DHCPv6 server in pfSense that makes it that way. (When I look in the pcap file with wireshark the Client Identifier for all 4 DHCPv6 packets show ...:00:11:22:33:44:55)
*) In particular I think of the MAC address part in the DUID that becomes 00:11:22:22(<--double 22!):33:44:55 in the pfSense GUI.

I'm not into this stuff so it is a bit confusing. I am not at all into how IA-NA and IA-PD are calculated/determined, but some of it seems to happen inside of the DHCPv6 service in pfSense before it gets into the leases file.

4. The dhcpd6.leases file content for IA-NA and IA-PD also looks different compared to the description of this bug report, because now it does not only contain octal numbers. I guess that must be because of the new MAC addresses I set before I started the capture.

#5 - 01/28/2015 10:58 PM - Chris Buechler

- Target version changed from 2.2.1 to 2.2.2

#6 - 04/03/2015 06:13 PM - Chris Buechler

- Target version changed from 2.2.2 to 2.2.3

#7 - 06/01/2015 06:30 PM - Chris Buechler

- Target version changed from 2.2.3 to 2.3

#8 - 01/15/2016 11:11 PM - Chris Buechler

- Priority changed from Normal to Low

- Target version deleted (2.3)

Not seeing an easy way to match up in this circumstance that won't potentially over-match or cause other problems. Rare situation that's never worked and only one person's ever noticed.

#9 - 01/18/2016 02:18 PM - Anders Lind

Thank you for the update Chris.
I hope it is not too late to include my patch for pfSense 2.3 regarding prefixes.php ?

It is about https://github.com/pfsense/pfsense/pull/2470

I did some investigation over the last weeks and made a patch addressing the mentioned issue + another that occurs in 1 out of 256 times per octet in the IAID+DUID string (actually one of the non-octal encoded octets in the leases file: \n)
so that routing and IPv6 internet connectivity now works!
Not addressed by my patch (because it is unrelated and problem lies somewhere else in pfSense) is some GUI related quirks that I describe in the end. The file that likely needs to be fixed is: https://github.com/pfsense/pfsense/blob/master/src/usr/local/www/status_dhcpv6_leases.php

To understand how prefixes.php uses the ISC DHCPv6 lease file and how it works I made a converter/decoder to interpret the lease content:
https://anderslind.dk/isc-dhcpdv6-lease-decoder

I flashed my sub-router (which is a home router) D-Link DIR-860L A1 (originally running firmware 1.08) with DD-WRT instead and have been using that for quite some time.
I have made tests with different client dhcp6 configurations using DD-WRT and pfSense to reproduce and test the bugs mentioned above and through these tests I noticed those other (GUI related) bugs and made tests for those as well.

I also read https://redmine.pfsense.org/issues/2347 e.g. about the two options/paths Seth mentions regarding adding routes to delegated IPv6 prefixes.

So I try to wrap it up now.

It seems it is ok that the IAID of the IA_NA != that of the IAID of the IA_PD according to RFC 3633 https://tools.ietf.org/html/rfc3633#section-6
Please see section 6 with respect to paragraph 2 (Important):

> An IA_PD is different from an IA, in that it does not need to be
> associated with exactly one interface. One IA_PD can be associated
> with the requesting router, with a set of interfaces or with exactly
> one interface. A requesting router must create at least one distinct
> IA_PD. It may associate a distinct IA_PD with each of its downstream
> network interfaces and use that IA_PD to obtain a prefix for that
> interface from the delegating router.

(compared to IAID description for IA_NA found in RFC 3315 https://tools.ietf.org/html/rfc3315#section-22.4 :

> IAID
> The unique identifier for this IA_NA; the
> IAID must be unique among the identifiers for
> all of this client's IA_NAs. The number
> space for IA_NA IAIDs is separate from the
> number space for IA_TA IAIDs.)

as well as sections:
section 7 with respect to paragraph 1 and 2 https://tools.ietf.org/html/rfc3633#section-7
section 8 with respect to paragraph 1 and 2 https://tools.ietf.org/html/rfc3633#section-8

Basically the patch I made is about:
1. Handling of IA_NA and IA_PD strings (that contain IAID+DUID content) using only the DUID part.
2. Fixing regular expression matching with respect to the IAID+DUID string regarding the legal \" substring (used in ISC DHCPv6 leases).
3. Checking the $duid variable before use. Default case for $type in the switch case statement.

Regarding 1. Especially section 8 paragraph 2 of RFC 3633 says SHOULD (https://tools.ietf.org/html/rfc2119) regarding:

> When a requesting router sends a DHCP message, it SHOULD be sent on
> the interface associated with the upstream router (ISP network). The
> upstream interface is typically determined by configuration. This
> rule applies even in the case where a separate IA_PD is used for each
> downstream interface.

, so it should be quite safe to ignore the IAID part! Maybe for relay agents something extra is needed https://tools.ietf.org/html/rfc3633#section-14 and Errata ID: 3736 (in the bottom) https://www.rfc-editor.org/errata_search.php?rfc=3633 , but I assume this is quite out of scope. (It seems to be undefined how this should work.)

Regarding 2. this means the longer the IAID+DUID string is the more likely it is that it includes a \" and that the regular expression matches too soon, because it is non-greedy! Please read the patch for more comments about this.

Tests

I have tested the patch under different circumstances:
In all the cases below the route was created from pfSense to the sub-router and the sub-router had then working IPv6 connectivity to the Internet:

```
Test A:
Conditions:
IAID=normal (including neither " nor \ somewhere)
DUID=normal mac address (including neither " nor \ somewhere)
ia-na IAID(+DUID) == ia-pd IAID(+DUID)
Result: Works!

Test B:
Conditions:
IAID=normal (including neither " nor \ somewhere)
DUID=normal mac address (including neither " nor \ somewhere)
ia-na IAID(+DUID) != ia-pd IAID(+DUID)
Result: Works!
```

In the following tests the pfSense GUI "Status / DHCPv6 leases" does not present either the right IAID or the right DUID whenever a " or a \ is used (but routing is established and works!):
(d=decimal, h=hex)

```
Test C:
DUID=mac address as 00 11 22 33 44 55 (notice 22h==34d==")
IAID=normal (including neither " nor \ somewhere)
ia-na IAID(+DUID) != ia-pd IAID(+DUID)
Result: Works, but the pfSense GUI presents the DUID as:
00:03:00:01:00:11:22:22:33:44:55
when in reality it is:
00:03:00:01:00:11:22:33:44:55
It is both a problem in Leases and Delegated Prefixes in the GUI that the DUID is presented as 22:22 where it should just be 22.

Test D:
IAID=34d  (") Tested with IAID of IA-NA
DUID=normal mac address (including neither " nor \ somewhere)
ia-na IAID(+DUID) != ia-pd IAID(+DUID)
Result: Works, but the pfSense GUI shows that IAID is 8738 which in hex is 2222 == "" (two apostrophes)

Test E:
IAID=34d  (") Tested with IAID of IA-PD
DUID=normal mac address (including neither " nor \ somewhere)
ia-na IAID(+DUID) != ia-pd IAID(+DUID)
```

```
Result: Works, but the pfSense GUI shows that IAID is 8738 which in hex is 2222 == "" (two apostrophes)

Test F:
DUID=mac address as 00 11 5C 33 44 55 (notice 5Ch==92d==\)
IAID=normal (including neither " nor \ somewhere)
ia-na IAID(+DUID) != ia-pd IAID(+DUID)
Result: Works, but the pfSense GUI presents the DUID as:
00:03:00:01:00:11:5c:03
when in reality it is:
00:03:00:01:00:11:5c:33:44:55

Test G:
IAID=92d  (\) Tested with IAID of IA-NA
DUID=normal mac address (including neither " nor \ somewhere)
ia-na IAID(+DUID) != ia-pd IAID(+DUID)
Result: Works, but the pfSense GUI presents the IAID as 23644 (5C 5C in hex). Also Wireshark confirms that bot
h sub router and pfSense talk about one 5C (and not two 5Cs).

Test H:
IAID=92d  (\) Tested with IAID of IA-PD
DUID=normal mac address (including neither " nor \ somewhere)
ia-na IAID(+DUID) != ia-pd IAID(+DUID)
Result: Works, but the pfSense GUI presents the IAID as 23644 (5C 5C in hex).
```

The GUI also needs a little fixing with regards to column headings and content that are not aligned on the "Status / DHCPv6 leases" page. At least in the alpha release I have:
2.3-ALPHA (amd64)
built on Sat Jan 02 09:55:38 CST 2016

Note:
It is my plan later to make a live test with all the CPEs that are present in my local non-profit housing association and see how this patch differs to the old version of prefixes.php to get an overview about how many CPEs (Ethernet SOHO routers) actually ask for a prefix and get routes created. But that will first happen in the coming months.

#10 - 01/21/2016 11:32 PM - Jim Thompson

- Target version set to 2.3

#11 - 01/21/2016 11:32 PM - Jim Thompson

please evaluate https://github.com/pfsense/pfsense/pull/2470

#12 - 01/22/2016 12:00 AM - Anders Lind

- Status changed from Confirmed to Feedback

- % Done changed from 0 to 100

Applied in changeset 2caea6a2cca2cbc37f1d99ebd8e5ea8277661d77.

#13 - 01/22/2016 12:05 AM - Chris Buechler

thanks Anders, appreciate the leg work on that. Looks to test out fine, will leave to feedback for additional testing.

#14 - 01/23/2016 05:37 AM - Chris Buechler

- Status changed from Feedback to Resolved

Works.

Thanks Anders!

Files

packetcapture.cap	5.08 KB	01/20/2015	Anders Lind
Status DHCPv6 leases.png	165 KB	01/20/2015	Anders Lind
dhcpd6.leases	892 Bytes	01/20/2015	Anders Lind