

## pfSense - Bug #5993

### dhcp6c not started until an RA received

03/14/2016 05:20 PM - Richard Patterson

<b>Status:</b>	Resolved	<b>Start date:</b>	03/14/2016
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Renato Botelho	<b>% Done:</b>	0%
<b>Category:</b>	DHCP6	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	2.4.0		
<b>Affected Version:</b>	All	<b>Affected Architecture:</b>	All

#### Description

On a residential IPoE (not PPPoE) broadband line, a BNG will often require a subscriber to send a DHCPv6 SOLICIT in order for it to be authenticated on to the ISP's network.

Currently pfsense uses rtsold which waits for an RA to be received before calling the rtsold\_<iface>\_script.sh which in turn triggers dhcp6c. As the ISP's BNG never sends an RA, this causes a stalemate with the ISP and no IPv6 PD is delegated.

#### History

##### #1 - 03/16/2016 08:46 AM - Richard Patterson

- File *interfaces.inc.patch* added

Attached rudimentary patch for *interfaces.inc* proves that starting *dhcp6c* prior to *rtsold* is sufficient to trigger the ISP to authenticate the client and start sending RAs.

It however still allows the *rtsold\_<iface>\_script.sh* to kill and restart (superfluously) *dhcp6c* once it receives the initial RA from the ISP.

See discussions here: [\[\[https://forum.pfsense.org/index.php?topic=102008.msg568937#msg568937\]\]](https://forum.pfsense.org/index.php?topic=102008.msg568937#msg568937)

Patch was against latest 2.3-BETA build.

##### #2 - 03/16/2016 09:52 AM - Martin Wasley

- File *interfaces.inc.patch* added

Richard's patch although allowed *dhcp6c* to run and get a prefix delegation has issues with losing that prefix shortly afterwards when the *RTSOLD* script is run.

The patch I have attached here launches the *dhcp6c* daemon at the end of the *ipv6* interface configure and stops *RTSOLD* from launching it. I have also cleaned up the killing of the *dhcp6c* process by calling a function in the same way as the *IPV4* client.

This is a fix for *IPoE*, for *PPoE* then it may be that *RTSOLD* needs to launch the *dhcp6c* daemon and not at the end of the configure, hence the comments I have put in the patch.

All that is needed is the addition of a GUI option for *IPoE* or *PPoE* and then a couple of *if..else* statements.

Also of note, I no longer see any *XID\_MISMATCH* errors.

**#3 - 03/16/2016 11:49 AM - Martin Wasley**

- File *interfaces\_2.3.inc.patch* added

Sorry just noticed patch I posted earlier is backwards!! Mea culpa!

**#4 - 03/17/2016 02:32 AM - Chris Buechler**

- Status changed from *New* to *Confirmed*

- Assignee set to *Chris Buechler*

- Target version set to *2.3.1*

There are a number of potential race conditions in the patches here, the one Martin noted among other possibilities. It would also introduce problems like noted in [#5999](#) in some edge cases. Given the status quo works for virtually everyone, and we're close to RC, I'll take a look at it for 2.3.1.

The patches here are probably OK for people who are in the specific circumstance noted. The edge case bugs it introduces probably won't affect most users.

**#5 - 03/18/2016 01:31 PM - Martin Wasley**

- File *interfaces.inc* added

- File *interfaces.php* added

Chris Buechler wrote:

There are a number of potential race conditions in the patches here, the one Martin noted among other possibilities. It would also introduce problems like noted in [#5999](#) in some edge cases. Given the status quo works for virtually everyone, and we're close to RC, I'll take a look at it for 2.3.1.

The patches here are probably OK for people who are in the specific circumstance noted. The edge case bugs it introduces probably won't affect most users.

Ah.. continuing on from message on the forum I now see why I could not add to the thread, I was not logged in. :)

**#6 - 03/18/2016 06:44 PM - Chris Buechler**

- File *deleted (interfaces.inc.patch)*

**#7 - 03/18/2016 06:44 PM - Chris Buechler**

- File *deleted (interfaces\_2.3.inc.patch)*

**#8 - 03/18/2016 06:48 PM - Chris Buechler**

- File *deleted (interfaces.php)*

**#9 - 03/18/2016 06:48 PM - Chris Buechler**

- File deleted (interfaces.inc)

**#10 - 03/18/2016 06:49 PM - Chris Buechler**

- File redmine5993.patch added

attaching a patch version of the full files Martin posted.

**#11 - 03/19/2016 03:46 AM - Martin Wasley**

Just a note to anyone who uses the patches, the patch does two files, interfaces.inc and interfaces.php. There is a new option in WAN setup which allows the selection of the dhcp6c client to run before or after RA, I have called these ipoe or pppoe in the IPV6 DHCP section. Selecting pppoe will cause dhcp6c and rtsold to run in the original way, selecting ipoe will reverse this and dhcp6c will NOT run when the rtsold script is executed. By default ipoe is selected, so if you want it to run in the original way you will need to select that in the WAN setup.

**#12 - 03/25/2016 09:07 AM - Martin Wasley**

- File interfaces.inc.patch added

The interfaces.inc.patch attached was generated against 2.3.b.20160325.07\*\*, I have applied it to my system and all is working. The interfaces.php.patch part of the redmine5993.patch still works, so run the redmine patch first, then the patch attached against interfaces.inc.

**#13 - 04/06/2016 08:21 AM - Martin Wasley**

- File interfaces.inc.patch added

- File interfaces.php.patch added

Attached are patches generated against 2.3.r.20160405.2024

interfaces.inc changes made to improve the logic around the old rtsold creation.

interfaces.php changes made add extra information to option box.

**#14 - 04/15/2016 10:08 PM - Chris Buechler**

- Target version changed from 2.3.1 to 2.3.2

**#15 - 04/19/2016 08:26 PM - David Wood**

The issue raised in this bug should certainly be addressed, but I urge careful review of the commentary in Issue 1 of <https://forum.pfsense.org/index.php?topic=101967.msg584693#msg584693> (which relates to [#5621](#)) before any further changes are made to the code in this area. In particular, I urge careful review of the matters in the 'Outstanding issues requiring further work' section of that post.

Two dhcp6c processes running on the same interface almost invariably breaks IPv6 connectivity on that interface, so extreme care is needed, especially as previous experience has shown how fertile a ground this area is for timing issues. Checking for an existing dhcp6c process on the interface is not enough to guarantee that there is no possibility of two dhcp6c processes on that interface. If interface\_dhcp6c\_configure() has started rtsold -1 on an interface, that rtsold will trigger a new dhcp6c when the next RA is received.

I am certainly not defending the current code, which is ripe for improvement. Getting things from where they were in 2.2.4 to where they are today (by addressing [#5297](#) and [#5621](#)) has taken me probably 50-100 hours of development time, as well as the willingness of several forum members to test patches and offer their suggestions. As probably the last person to spend a lot of time with this code, I wanted to flag up how kludgy the current code is even after my improvements, also how fragile it often proves to be when changes are made.

It might be better to tidy up and refactor the existing code, then deal with the known problems, before adding any more features to code that has to handle many different configuration scenarios and can destroy IPv6 connectivity if it malfunctions.

Probably the key architectural issue in this area is that pfSense currently does not keep track of whether it believes dhcp6c or rtsold is supposedly running on an interface. There isn't even a single 'shut down dhcp6c' routine at present - instead there are calls in three separate files to posix\_kill() dhcp6c (and Martin's patch proposes adding yet another routine to kill dhcp6c). When I wrote the /usr/local/sbin/ppp-ipv6 script, I unset the nd6 ACCEPT\_RTADV setting (controlling the kernel's response to receiving an RA) partly to prevent action on a spurious RA received when the link was

supposed to be down, but mainly as a sort of lock, telling the link up routine that the link down routine had run and should have killed any existing dhcp6c process.

It might be better to start any further work in this area by making pfSense track the state of rtsold and dhcp6c on a per-interface basis (presumably by pid), giving these daemons the proper locking they should have by reworking /usr/local/sbin/ppp-ipv6, /etc/inc/interfaces.inc, /etc/rc.newwanip and /usr/local/www/interfaces.php. This would be a good foundation for addressing the issues I've highlighted in Issue 1 of that [forum post](#). Once this work is done, the code should be more robust and better prepared for further changes.

I'm fairly certain that Track Interface still does not do anything sensible on a change of delegated prefix(es) - that might be best addressed before adding in this feature, though a full reworking of the Track Interfaces functionality is a big job whereas the feature proposed here is a relatively minor enhancement that is crucial for IPv6 connectivity in a small but important subset of environments.

#### **Martin - some suggestions**

When submitting patches, it's best to submit patches in unified diff format, i.e. diff -u, as that reduces their version sensitivity and allows Redmine's preview functionality to work properly. It is best to include all patches in a single file.

It is even better if you [fork](#) the [GitHub repository](#) and work on a feature branch in your fork, as this opens up the full range of tools on GitHub and makes it trivial to submit a pull request to the pfSense developers.

#### **What should the flag be called?**

I dislike the proposed description of the flag in interfaces.php. This is not necessarily an IPoE versus PPPoE thing, though it is likely that PPPoE BNGs will send periodic RAs before receiving a DHCPv6 solicit - see [IR-187](#) section 5.2 ("In all cases, the BNG generates Router Advertisement messages toward the PPP peer."), 9.1 R-24, and 9.2 R-46 and R-57.

Maybe a clearer description is:

Do not wait for a RA [ ] (Required by some ISPs, especially those not using PPPoE)

and call the flag itself dhcp6withoutra. As the checkbox would be in the section relating to DHCP6, I don't see the need to add "before starting DHCP6" to its title.

#### **#16 - 04/20/2016 03:19 PM - Richard Patterson**

David Wood wrote:

#### **What should the flag be called?**

I dislike the proposed description of the flag in interfaces.php. This is not necessarily an IPoE versus PPPoE thing, though it is likely that PPPoE BNGs will send periodic RAs before receiving a DHCPv6 solicit - see [IR-187](#) section 5.2 ("In all cases, the BNG generates Router Advertisement messages toward the PPP peer."), 9.1 R-24, and 9.2 R-46 and R-57.

Maybe a clearer description is:

Do not wait for a RA [ ] (Required by some ISPs, especially those not using PPPoE)

and call the flag itself dhcp6withoutra. As the checkbox would be in the section relating to DHCP6, I don't see the need to add "before starting DHCP6" to its title.

I'd say not to bother having a flag/option and to just make this the default behaviour.

If a user has set DHCPv6 on the WAN, they want it to send a SOLICIT. I can't think of any possible downsides to not waiting for an RA.

The only reason it's waiting for an RA currently is because it's utilising rtsold to launch an external script to start dhcp6c. Which is odd, IMO, and I can only assume it's a side effect behaviour of having a userland RA/RS daemon.

I say odd, but the reasoning could make sense if it was waiting for the RA to see if the M/O option flags are set. However if a user has configured pfsense to use stateful DHCPv6, it should do so regardless.

#### #17 - 04/21/2016 03:19 AM - David Wood

Richard Patterson wrote:

David Wood wrote:

Maybe a clearer description is:

Do not wait for a RA [ ] (Required by some ISPs, especially those not using PPPoE)

and call the flag itself dhcp6withoutra. As the checkbox would be in the section relating to DHCP6, I don't see the need to add "before starting DHCP6" to its title.

I'd say not to bother having a flag/option and to just make this the default behaviour.

If a user has set DHCPv6 on the WAN, they want it to send a SOLICIT. I can't think of any possible downsides to not waiting for an RA.

The precautionary principle suggests that changing the current behaviour for the vast majority of users for whom it works is unwise, especially as this is an area that has thrown up timing issues in the past. The current behaviour has been in place since [29f2f07a66f880732ee6656eda745e7b108db7e4](https://github.com/pfsense/pfsense/commit/29f2f07a66f880732ee6656eda745e7b108db7e4), which was made before the release of pfSense 2.1 (the first version with IPv6 support).

If a flag and option are implemented, 'Do not wait for a RA' can eventually become the default if that proves better than the current 'RA before dhcp6' behaviour for the majority of users. A flag/option allows anyone who requires the current behaviour to opt out of the new default easily, also those for whom the current behaviour works can be left on that behaviour rather than being switched to the new default.

In the light of further reflection on the existing code (see below), I propose that a better name for the flag might be dhcp6beforera.

Richard Patterson wrote:

The only reason it's waiting for an RA currently is because it's utilising rtsold to launch an external script to start dhcp6c. Which is odd, IMO, and I can only assume it's a side effect behaviour of having a userland RA/RS daemon.

I say odd, but the reasoning could make sense if it was waiting for the RA to see if the M/O option flags are set. However if a user has configured pfsense to use stateful DHCPv6, it should do so regardless.

SLAAC is implemented entirely in the kernel, but FreeBSD has no capacity to solicit routers in the kernel - that's where rtsold comes in. pfSense needs to run rtsold in order to find the router on the interface.

I'm going to start some of the clean up work in a feature branch in my fork of the repository, with the intention of implementing this feature along the way. When I have anything to share, I'll post a link to the branch and details of how to patch the latest release.

**#18 - 04/21/2016 03:38 AM - Chris Buechler**

David Wood wrote:

The precautionary principle suggests that changing the current behaviour for the vast majority of users for whom it works is unwise, especially as this is an area that has thrown up timing issues in the past. The current behaviour has been in place since [29f2f07a66f880732ee6656eda745e7b108db7e4](https://github.com/pfsense/pfsense/commit/29f2f07a66f880732ee6656eda745e7b108db7e4), which was made before the release of pfSense 2.1 (the first version with IPv6 support).

If a flag and option are implemented, 'Do not wait for a RA' can eventually become the default if that proves better than the current 'RA before dhcp6' behaviour for the majority of users. A flag/option allows anyone who requires the current behaviour to opt out of the new default easily, also those for whom the current behaviour works can be left on that behaviour rather than being switched to the new default.

This is wise, indeed. That's the way I want to see this implemented, for exactly the reasons noted. If it works out fine, then maybe we change the default for new configs to use that option in the future.

I'm going to start some of the clean up work in a feature branch in my fork of the repository, with the intention of implementing this feature along the way. When I have anything to share, I'll post a link to the branch and details of how to patch the latest release.

That'd be great. It's more than I have time to get into at the moment and not significant enough to make the cut for a quick turnaround 2.3.1.

**#19 - 05/11/2016 06:00 AM - Martin Wasley**

Chris Buechler wrote:

David Wood wrote:

The precautionary principle suggests that changing the current behaviour for the vast majority of users for whom it works is unwise, especially as this is an area that has thrown up timing issues in the past. The current behaviour has been in place since [29f2f07a66f880732ee6656eda745e7b108db7e4](https://github.com/pfsense/pfsense/commit/29f2f07a66f880732ee6656eda745e7b108db7e4), which was made before the release of pfSense 2.1 (the first version with IPv6 support).

If a flag and option are implemented, 'Do not wait for a RA' can eventually become the default if that proves better than the current 'RA before dhcp6' behaviour for the majority of users. A flag/option allows anyone who requires the current behaviour to opt out of the new default easily, also those for whom the current behaviour works can be left on that behaviour rather than being switched to the new default.

This is wise, indeed. That's the way I want to see this implemented, for exactly the reasons noted. If it works out fine, then maybe we change the default for new configs to use that option in the future.

I'm going to start some of the clean up work in a feature branch in my fork of the repository, with the intention of implementing this feature along the way. When I have anything to share, I'll post a link to the branch and details of how to patch the latest release.

That'd be great. It's more than I have time to get into at the moment and not significant enough to make the cut for a quick turnaround 2.3.1.

Sorry guys, I have been away from home for a couple of weeks or so working in Dubai, and previous to that I was getting a project ready for that trip so I have not been following this thread.

The flag I implemented made it a default for me and a couple of others who were testing in order to make life easier. I had intended to reverse that flag meaning I and others needing that option would need to set it, I just never got around to it.

I am a little confused David, and it happens a lot these days... Are you saying that rtsold will launch dhcp6c even though I have specifically made it so that the script does not launch it?

The only issue I have had with this modification is that when doing a WAN down and then up again I do get multiple dhcp6c clients, even though I have scanned for all instances and killed them, that would be fine for a user with only one WAN interface but no good for multiple WAN, as you have said the current creation of dhcp6c\_({\$wanif}).pid falls over when dhcp6c gets launched again as that file is overwritten, I should add that this multiple dhcp6c issue only ever happens in this WAN down and back up situation.

Now, putting some debugging in there, because I also modified the dhcp6c client to have a 'no-release' option indicated to me that the function interface\_dhcpv6\_configure was being called multiple times as the command line was always the same, thus indicating it was being launched where I had put it, and not by rtsold as that had the no-release option omitted from the command, however even more bizarre is that immediately after launching dhcp6c I send a logger message marked 'mwttag "Starting dhcp6 client for interface wan({\$wanif} in IPoE mode)" and this only appears once, thus totally bemused I left it there and went and did some paper work! :)

David Wood wrote:

When submitting patches, it's best to submit patches in unified diff format, i.e. diff -u, as that reduces their version sensitivity and allows Redmine's preview functionality to work > properly. It is best to include all patches in a single file.

It is even better if you fork the GitHub repository and work on a feature branch in your fork, as this opens up the full range of tools on GitHub and makes it trivial to submit a pull > request to the pfSense developers.

Being my first trip into the world of GitHub I had a few problems getting windows and Github to behave, probably all down to me no doubt, but I'll endeavour to RTFM and follow your suggestions.

**#20 - 07/06/2016 08:59 AM - Martin Wasley**

I finally managed to get back to this after several weeks having to work for a living. The first thing I did was to update to the latest release.

Having done that, I decided to look deeper into the problem of multiple dhcp6c clients running... guess what, it's not happening now. I have a test pfSense running on a VM running the original 2.3 and it does happen on that, so something has changed that has fixed that issue.

I have updated interfaces.php and interfaces.inc with the modifications I use to make it work with my ISP is working perfectly now, no issues at all.

**#21 - 07/11/2016 09:41 PM - Kevin Morse**

Martin Wasley wrote:

I finally managed to get back to this after several weeks having to work for a living. The first thing I did was to update to the latest release.

Having done that, I decided to look deeper into the problem of multiple dhcp6c clients running... guess what, it's not happening now. I have a test pfSense running on a VM running the original 2.3 and it does happen on that, so something has changed that has fixed that issue.

I have updated interfaces.php and interfaces.inc with the modifications I use to make it work with my ISP is working perfectly now, no issues at all.

Hi Martin,

Do you mind posting patches based on the changes you've made? The most recent set of patches appear to be from April.

Thanks,

Kevin

**#22 - 07/17/2016 09:16 PM - Chris Buechler**

- Status changed from *Confirmed* to *Feedback*

- Assignee deleted (*Chris Buechler*)

- Target version changed from *2.3.2* to *2.4.0*

merged this for 2.4 as it needs more baking time in snapshots than we're going to have for 2.3.2.

**#23 - 07/19/2016 08:51 PM - Daryl Morse**

Chris Buechler wrote:

merged this for 2.4 as it needs more baking time in snapshots than we're going to have for 2.3.2.

I just downloaded the latest 2.3.2 snapshot, pfSense-CE-2.3.2-DEVELOPMENT-amd64-20160718-2327.iso and installed the patch. My settings are request only prefix, /56, send hint and do not wait. I tried rebooting both pfsense and the client pc (running windows 10), but it made no difference. I tried applying the patch to both the stable and developmental snapshots branches, but neither worked. In both cases, by manually starting dhcp6c (/usr/local/sbin/dhcp6c -c /var/etc/dhcp6c\_wan.conf -p /var/run/dhcp6c\_hn0.pid hn0), I was able to get an ipv6 address on the pc. The dhcp6 gateway stays in pending status, never comes online.

Let me know if I should be using some other snapshot or other methodology to test or if I should post logs.

**#24 - 07/25/2016 08:26 PM - Daryl Morse**

Chris Buechler wrote:

merged this for 2.4 as it needs more baking time in snapshots than we're going to have for 2.3.2.

I posted an update about this patch (now two patches) in this thread: <https://forum.pfsense.org/index.php?topic=114511.0>

**#25 - 08/01/2016 12:22 PM - Daryl Morse**

Daryl Morse wrote:

Chris Buechler wrote:

merged this for 2.4 as it needs more baking time in snapshots than we're going to have for 2.3.2.

I posted an update about this patch (now two patches) in this thread: <https://forum.pfsense.org/index.php?topic=114511.0>

I've been using these patches for a week now. I was originally running with a single windows 10 client, but I created some additional vms so there are now 4 clients. Another user from my ISP is also testing it. As far as I know, it's working properly for him. Hopefully tomorrow, I'll have time to switch my LAN over to use this patch, which will be a bigger test of 20-30 \*nix, windows and android clients.

**#26 - 08/05/2016 12:59 AM - Daryl Morse**

I still haven't had a chance to switch my LAN over to this software, but I'm aware of three other Telus users who are running 2.3.2 with the above two patches. In all three cases, as far as I know, they have multiple hosts of various different types and they not reported any issues.

I'm currently testing the latest development snapshot (as of the time of this post) with the latest patch (Corrections and additions to dhcp6 before RA

[#3087](#)). This patch addresses the issue of there being multiple instances of dhclient and dhcp6c. However, there are some issues.

The first issue pertains to the status of WAN\_DHCP6. When the system initially boots, the status is initially reported as unknown with the metrics shown as pending. However, the gateway actually is working and all 6 dhcp-related processes are started (`ps auxw | grep dhcp`) and working properly (as evidenced by dhcp4 and dhcp6 leases being granted). After releasing the WAN interface, WAN\_DHCP goes offline, WAN\_DHCP6 stays unknown. After renewing the WAN interface, both gateways go online. After releasing the WAN interface again, WAN\_DHCP goes offline, WAN\_DHCP6 stays online. So, the dhclient and dhcp6c processes are starting up and shutting down properly, but the status of WAN\_DHCP6 is not being reported properly, even though the processes appear to be working properly.

There are a couple of other issues that I don't think are related to this patch, at least not directly, because I also noticed them without the latest patch installed.

The first issue is the management of dhcp6 server processes (`dhcpcd -6` and `dhcpleases6`). Every time the WAN interface is released and renewed another pair of processes appears. The orphan processes don't seem to prevent dhcp6 from working. I've found that by stopping the service, manually killing the orphans and restarting the service, it starts to work.

I've also noticed that dhcp and dhcp6 lease status doesn't always report properly, even though leases are being granted (I can tell that by using `ipconfig /all` on the client).

Although I have no way to test the above without using the dhcp before RA feature, I wouldn't be surprised if the same issues are still there.

I hope that helps. If any other info or log files are required, let me know what you need.

#### **#27 - 08/06/2016 12:32 PM - Daryl Morse**

Daryl Morse wrote:

I'm currently testing the latest development snapshot (as of the time of this post) with the latest patch (Corrections and additions to dhcp6 before RA [#3087](#)). This patch addresses the issue of there being multiple instances of dhclient and dhcp6c. However, there are some issues.

The first issue pertains to the status of WAN\_DHCP6. When the system initially boots, the status is initially reported as unknown with the metrics shown as pending. However, the gateway actually is working and all 6 dhcp-related processes are started (`ps auxw | grep dhcp`) and working properly (as evidenced by dhcp4 and dhcp6 leases being granted). After releasing the WAN interface, WAN\_DHCP goes offline, WAN\_DHCP6 stays unknown. After renewing the WAN interface, both gateways go online. After releasing the WAN interface again, WAN\_DHCP goes offline, WAN\_DHCP6 stays online. So, the dhclient and dhcp6c processes are starting up and shutting down properly, but the status of WAN\_DHCP6 is not being reported properly, even though the processes appear to be working properly.

Further to what I quoted above, I found that by restarting the dpinger service on the dashboard, the WAN\_DHCP6 gateway status correctly shows online. Based on that, this patch seems to be working rather well.

**#28 - 08/07/2016 06:07 PM - Daryl Morse**

I'm currently testing the latest development snapshot (as of the time of this post) with the latest patch from (DHCP6 Before RA. Additions and ammendments [#3092](#)). The system behaviour is similar to what I reported above. When pfsense initially boots, the status of WAN\_DHCP6 is reported as unknown. After restarting the dpinger service, the status is reported as online.

When I took a look at the code, I noticed a typo on line 3971: `$rtsoldscript .= "/usr/bin/logger -t rtsold \"Recieved RA specifying route \\\$2 for interface {\$interface}({\$wanif})\"\\n";` "Recieved" is spelled incorrectly.

**#29 - 08/07/2016 10:57 PM - Phillip Davis**

The text typo was in existing code, so I made a separate pull request to tidy that up:

<https://github.com/pfsense/pfsense/pull/3094>

**#30 - 08/08/2016 03:19 PM - Daryl Morse**

Phillip Davis wrote:

The text typo was in existing code, so I made a separate pull request to tidy that up:

<https://github.com/pfsense/pfsense/pull/3094>

Thanks, Phil. I figured the typo was a legacy.

**#31 - 08/08/2016 10:31 PM - Daryl Morse**

Phillip Davis wrote:

The text typo was in existing code, so I made a separate pull request to tidy that up:

<https://github.com/pfsense/pfsense/pull/3094>

I tried to apply your patch, but it won't go in with PR [#3092](#) installed, nor the other way around. That's unfortunate.

**#32 - 08/08/2016 10:37 PM - Daryl Morse**

I looked into the gateway status issue. After pfsense boots with PR [#3092](#) installed, the status of WAN\_DHCP6 is unknown. There is only one instance of dpinger running, and it's for WAN\_DHCP. By restarting the dpinger service, both instances of dpinger are running and the status is correctly shown as online for both gateways.

**#33 - 08/13/2016 01:38 AM - Daryl Morse**

I was holding off on upgrading to the latest snapshot because PR 3092 wouldn't install. However, I noticed today that Jorge / NewEraCracker created PR 3102, so I backed out of PR 3092, then installed PR 3094 and PR 3102. It worked, so I upgraded to the latest snapshot and installed only PR 3102. It looks good, although it has exactly the same problem as I reported above, which is that WAN\_DHCP6 does not show online status after

booting until dpinger is restarted. I looked in the log and there is nothing obvious, aside from no dpinger messages for WAN\_DHCP6 until after I restarted the service.

PS. I'm not sure if Jorge regularly checks this thread. I don't have an account on github. I'd appreciate if someone would ask him to read this thread. I'm bimmerdriver in the forum in case he wants to take this offline.

**#34 - 08/13/2016 04:42 PM - Jorge M. Oliveira**

Copying the info I shared on my reply to your PM on the forum.

From my understanding, the whole point to setting dhcp6c without RA, is so a link-local gateway can be used with another IPv6 given via DHCP. Thereby I believe that doing the following will sort your dpinger issues:

Go to "System > Routing > Gateways", edit the WAN\_DHCP6 gateway, and set the following advanced option "Use non-local gateway through interface specific route", save & apply.

Regards,  
Jorge M. Oliveira

**#35 - 08/13/2016 07:22 PM - Daryl Morse**

From my reply to your PM, based on a discussion with an engineer at my ISP, my understanding is the following:

The reason for this setting being required is because some ISPs decided to ignore RS until after DHCP6 solicit / advertise / request / reply (i.e., until after a prefix has been delegated). Only when a prefix has been delegated will their edge router respond to RS. It will probably make you laugh that my ISP claims the reason for this is "security". I guess you can call this "security by non-interoperability".

The standard TR-124 issue 4 describes router solicitation / dhcp6 prefix delegation process with a chart that clearly shows the RS/RA can take place concurrently with dhcp6 prefix delegation, however, some ISPs like mine have configured their edge routers, so pfsense without the patch will never get a prefix, because it's waiting forever for the RA, which will never come. It's a deadlock.

**#36 - 08/13/2016 07:46 PM - Jorge M. Oliveira**

I've updated my PR with another commit (almost the same I sent you via PM a few hours ago):

<https://github.com/pfsense/pfsense/pull/3102>

Needs testing. It now either kills and respawns dhcp6c (dhcp6withoutra == false) or it fires the dhcp6c\_script thereby calling /etc/rc.newwanipv6 (dhcp6withoutra == true). In either cases, system gets notified about the changes and everything is refreshed accordingly.

PS: This change has been tested on a small VM setup, it just needs testing in a physical environment to be sure it works well.

**#37 - 08/14/2016 01:19 AM - Daryl Morse**

I backed out of the previous changes (PR plus edits) and installed the updated PR. I tested it on a hyper-v server with 4 virtual windows 10 clients. It starts up properly (status and required processes). I'll beat on it more tomorrow.

**#38 - 08/14/2016 12:21 PM - Daryl Morse**

I spent a while doing some testing with pfsense and 4 clients. Your latest fix definitely seems to have fixed the problem I reported above, that the ipv6 gateway status does not go online after booting, but things go downhill from there. If I release the interface, the ipv4 gateway goes offline, but the ipv6 gateway status reports that it's online, no packet loss or anything. A while later, sometimes dpinger goes offline. Other times radvd and unbound go offline. If I renew the interface, the gateway statuses are both online, but sometimes I have to manually start radvd, which seems to bring back unbound. After this, I noticed there can be a duplicate dhcp6c process. I've also noticed when releasing and renewing the interface as well as restarting the dhcp service that in some cases, multiple dhcpd -6 and dhcpleases6 processes are started. If you would like to see this, I can show you using teamviewer.

**#39 - 08/15/2016 04:31 PM - Daryl Morse**

A lot of progress made on this bug and other issues. Currently, I'm running the latest snapshot with the following PRs installed: 3102/1, 3102/2, 3103, 3105, 3106 and 3107. The system boots up reliably with all dhcp-related processes running. This really good news and I appreciate the efforts of Martin and Jorge for getting it to this point.

The only remaining issue with respect to this specific bug is what I described above, WRT the status of WAN\_DHCP6 when the WAN interface is released. WAN\_DHCP goes offline right away, but WAN\_DHCP6 stays online indefinitely. In this state, RTT and RTTsd and packet loss update as if it's still running. Yesterday, I left the system in this state, came back a couple hours later and both gateways were in unknown state. If I restart dpinger, both gateways show unknown state. When the WAN interface is renewed, both gateways come online.

There are a couple of other things pertaining to dhcp6, but I'll report them in a separate bug.

**#40 - 08/15/2016 04:49 PM - Daryl Morse**

Correction:

I did another test. After around 30 minutes after releasing the WAN interface, both gateways were offline. They both came back online a unassisted after renewing the interface.

**#41 - 10/13/2016 03:51 PM - Daryl Morse**

Updating this issue based on 2.4 development snapshot.

The dhcp6 before RA feature has been working perfectly since the previous update. However, the behavior of WAN\_DHCP6 is unchanged. If the interfaces are released, WAN\_DHCP6 stays online for an extended period of time before finally going offline. Since not everyone can test this behavior with dhcp6 before RA set, let me know when there is a fix and I'll test it.

**#42 - 11/30/2016 02:45 AM - Martin Wasley**

Whilst having a look at another issue, the fabled no release on dhcp6c option, I noticed on WAN intergace startup that the dhcp6c\_\*\_script.sh runs twice. This is because it is still part of the rtsoldscript and that gets launched on receipt of an RA, but also dhcp6c also launches the dhcp6c\_\*\_script. I have commented out the dhcp6.\*.script.sh line that is used when dhcp6withoutra is selected and that solves the problem and I cannot see any side effects or any issues with my system.

Comments please.

Also, I am going to try and modify dhcp6c to not run the script if the PD has not changed on a refresh, this is pointless as far as I can see.

**#43 - 11/30/2016 03:53 AM - Jim Thompson**

- Assignee set to Jim Pingle

JimP, please look at the last entry here.

**#44 - 11/30/2016 08:27 AM - Jim Pingle**

- Assignee changed from Jim Pingle to Jim Thompson

I can see why it would end up being called twice since in certain combinations of configurations the script would end up in both the dhcp6c config and in the rtsold script. The specific logic behind the changes introduced in the patch isn't documented in the code, so it isn't clear if that's expected or required or a side effect of the settings on the interface.

If it works better without that other call, it's tempting to remove it, but I'd prefer to see some input from the PR's author @JorgeOliveira on what the intended effect was.

**#45 - 01/04/2017 02:41 AM - Martin Wasley**

The dhcpc before RA was originally my fix for an issue we have with Sky ISP in the U.K. I got very busy with work and Jorge came in and tidied it up, but the basic operation remains the same. I had looked at this before I started on the no-release and DUID additions but they took priority for me as the DUID and no-release were causing more problems. I had already removed this 'dual' call in my test device with no ill effects, and then had forgotten about it. I'll go and have a look at it again and see if I can break anything.

**#46 - 01/04/2017 10:19 AM - Martin Wasley**

OK, had a look around that bit of code. This is what I have found:

1. RTSOLD still launches multiple dhcp6c clients

I have re-written the rtsold script creation and added a lock file creation and check, it cannot launch dhcp6c if the lock file is present, the lock file is cleared in the same call to kill dhcp6c.

2. dhcp6c\*.pid still present after dhcp6c has exited.

Found that using kill -9 was causing this, changed it to kill -15 and a clean exit so the pid file is removed.

3. Very untidy around the bottom end of dhcp6c\_configure routine.

Moved the dhcpwithoutra launch to a separate routine and called that.

4. Removed the extra dhcp6c\_\$interface\_script call as it is not needed, left over from when I originally created the quick dirty fix for issues with ISP's wanting dhcp before RA.

Daryl can you try the patch I've done and let me know if that has solved your issues.

**#47 - 01/07/2017 01:20 AM - Martin Wasley**

The removal of the extra `dhcp6c_interface_script` call does cause a problem for some, those who use `dhcp6cwithoutra` and whose ISP's BNG is extremely slow in responding to the RS, in Daryl's case it was over four seconds after receiving the DHCP6 response that the RA was received, thus the script had run but there was no route so `dpinger` et al fails. This is an issue that only applies to `dhcp6cwithoutra`. Until a better solution can be found then leaving it to run twice is the only option.

**#48 - 01/10/2017 05:22 AM - Martin Wasley**

OK, it seems we have a solution. It involves a change to `dhcp6c`, another new flag is added!

The flag, currently 'x' prevents `dhcp6c` from running the script on receipt of 'REPLY' message. As 'REPLY' message should only ever be received once while the `dhcp6c` client is running, this solves the problem of the double run of `dhcp6c_wan_script`, as it is run by `RTSOLD` on receipt of the RA. This only applies to users running `dhcp6cwithoutRA`, so the flag is only used in that instance. The lock file creation I have added to `RTSOLD` appears to have stopped the launch of multiple `dhcp6c` clients, it appears this was being caused by multiple RA's being received in rapid succession.

Comments please, can anyone see any issues I haven't foreseen?

If not, then I and others will continue to test this for a few days more then I will issue PR's for `dhcp6c` and the changes I have made to `interfaces.inc`, then hopefully it should put this one to bed.

**#49 - 01/10/2017 06:30 AM - Martin Wasley**

Richard Patterson asked me by email to explain in more detail why I want to make these changes, here is my email to him which makes it easier to understand why it's needed.

pfSense normally sends out an RS, on getting the RA it then launches `dhcp6c`, `dhcp6c` runs the interface configure script on receipt of the Reply to the REQUEST response, having got the IA, PD r both.

With `dhcp6c` before RA, we stopped the `RTSOLD` script from launching `dhcp6c` and ran it before `RTSOLD`, the problem there is that although it works, the WAN configure script then gets run by both `dhcp6c` and `RTSOLD`. O.K. You say, then stop `RTSOLD` from running the WAN configure, that works fine with Sky, as the RA from Sky comes in very quickly so it arrives really before `dhcp6c` has run the WAN configure script. Now this is where there is a problem, we have found at least one ISP that delays it's RS response by several seconds, I don't know why, security is a possibility but the result is that the script has completed running before the RA has been received, the parameters from RA are therefore not pushed to the system and we end up with problems. Hence my idea of stopping `dhcp6c` from running the script on first receipt of a reply to a REQUEST response and letting `RTSOLD` launch it as normal.

<<<<

What I need to do is to change the logging messages coming from `DHCP6C` to say REPLY and not REQUEST.

**#50 - 01/17/2017 07:21 AM - Martin Wasley**

PR issued, should be the end of this one.

**#51 - 01/19/2017 11:55 AM - Martin Wasley**

As they say, this has been an experience.

This whole 5993 started because we needed dhcp6 before RA, and a quick fix was put in place. The last eight months have been spent trying to fix issues with that workaround, in other words create workarounds for the workaround.

I went back to square one and looked at it again, knowing much more than I did then. The result is that the solution was actually quite simple.

The default situation is that RTSOLD is launched, gets an RA and launches dhcp6c. The dhcp6c\_wan\_script runs rc.newwanipv6 on receipt of a response.

dhcp6withoutra requires things the other way around, we need to launch dhcp6c, get a response and then launch RTSOLD. So that is what I've done. Using the dhcpwithoutra flag it now creates a changed dhcp6c\_\*\_conf and replaces the existing script with a new script dhcp6c\_\*\_dhcp6withoutra\_script.sh which in turn launches RTSOLD.

dhcp6c having completed and set what it needs to then runs RTSOLD, RS gets sent and the RA to that then triggers RTSOLD to run the dhcp6c\_wan\_script which in turn runs rc.newwanipv6. When a renew event happens dhcp6c will again run RTSOLD and so on and so forth.

The dhcp6withoutRA is running on my live system and the default is running on my test system. Both are behaving.

Thoughts please.

**#52 - 01/20/2017 12:52 PM - Martin Wasley**

PR [#3410](#)

**#53 - 01/20/2017 02:42 PM - Daryl Morse**

Martin Wasley wrote:

PR [#3410](#)

I've been testing this PR since this morning. I've found that pfsense reliably acquires a prefix and ipv6 connectivity when booting, disabling / enabling the WAN interface and releasing / renewing the WAN interface. I've been monitoring the behaviour using wireshark. I appreciate all of effort of Martin and others who have worked on a fix to this bug. I hope this fix will be reviewed and tested by others so it can be included in the next release of pfsense. If anyone would like to see any logs or packet captures, don't hesitate to ask.

**#54 - 01/21/2017 12:03 PM - Daryl Morse**

I installed the additional patch that Martin provided to address the request for changes. I've tested both patches together including numerous reboots, wan interface disable / enable and wan interface release / renew, also with and without the flag to not allow dhcp release. I've been watching the icmp rs/ra and dhcpv6 packets using wireshark.

Martin, your fix is solid. It worked perfectly every time. Thank you very much for your efforts.

Please close this bug.

**#55 - 01/23/2017 12:29 AM - J L**

Daryl Morse wrote:

I installed the additional patch that Martin provided to address the request for changes. I've tested both patches together including numerous reboots, wan interface disable / enable and wan interface release / renew, also with and without the flag to not allow dhcp release. I've been watching the icmp rs/ra and dhcpv6 packets using wireshark.

Martin, your fix is solid. It worked perfectly every time. Thank you very much for your efforts.

Please close this bug.

Which patches exactly? Latest 2.4.0 snapshot I assume?

Thanks

**#56 - 01/23/2017 12:53 AM - Daryl Morse**

J L wrote:

Daryl Morse wrote:

I installed the additional patch that Martin provided to address the request for changes. I've tested both patches together including numerous reboots, wan interface disable / enable and wan interface release / renew, also with and without the flag to not allow dhcp release. I've been watching the icmp rs/ra and dhcpv6 packets using wireshark.

Martin, your fix is solid. It worked perfectly every time. Thank you very much for your efforts.

Please close this bug.

Which patches exactly? Latest 2.4.0 snapshot I assume?

Yes, I'm using the latest 2.4 beta snapshot.

If you follow the link to PR 3410 (<https://github.com/pfsense/pfsense/pull/3410>), there are two patches:

<https://github.com/pfsense/pfsense/pull/3410/commits/cdb6c8ac8e65f98a2ac0fa469c963c055a5c522d>  
<https://github.com/pfsense/pfsense/pull/3410/commits/5c803c910dd6c16e5563dbf13b0cbbbabe118e33>

#### #57 - 01/23/2017 02:20 AM - J L

Thanks.

I'm on the 2.4 snapshot and upon reboot I have to disable and enable the DHCPv6 server for it to properly delgate IPv6 addresses on my LAN. Annoying!

I tried with both patches and it happens still.

And I'm on Telus, hence the need for no RA.

Maybe I'll try reinstalling pfsense and starting fresh with these patches, or wait for 2.4 proper to be released.

Here's the of

```
ps -auxw |grep dhcp |grep -v grep
```

after startup with DHCPv6 not working.

```
_dhcp 11630 0.0 0.1 10496 2404 - Ss 02:02 0:00.00 dhclient: em4 (dhclient)
root 12955 0.0 0.1 8348 2264 - Is 02:02 0:00.00 /usr/local/sbin/dhcp6c -d -c /var/etc/dhcp6c_
wan.conf -p /var/run/
root 21934 0.0 0.1 8204 2084 - Ss 02:02 0:00.01 /usr/local/sbin/dhcpleases -l /var/dhcpd/var/
db/dhcpd.leases -d la
dhcpd 31680 0.0 0.4 24852 13908 - Ss 02:02 0:00.00 /usr/local/sbin/dhcpd -user dhcpd -group _dhcp
p -chroot /var/dhcpd
root 87637 0.0 0.1 10440 2520 - Ss 02:02 0:00.01 /usr/sbin/syslogd -s -c -c -l /var/dhcpd/var/
run/log -P /var/run/s
```

After disabling and enabling DHCPv6

```
_dhcp 11630 0.0 0.1 10496 2404 - Ss 02:02 0:00.00 dhclient: em4 (dhclient)
root 12955 0.0 0.1 8348 2264 - Is 02:02 0:00.00 /usr/local/sbin/dhcp6c -d -c /var/etc/dhcp6c_
wan.conf -p /var/run/
root 21934 0.0 0.1 8204 2084 - Ss 02:02 0:00.01 /usr/local/sbin/dhcpleases -l /var/dhcpd/var/
db/dhcpd.leases -d la
dhcpd 86906 0.0 0.4 24852 13908 - Ss 02:03 0:00.00 /usr/local/sbin/dhcpd -user dhcpd -group _dhp
p -chroot /var/dhcpd
dhcpd 87057 0.0 0.3 20756 11180 - Ss 02:03 0:00.00 /usr/local/sbin/dhcpd -6 -user dhcpd -group _
dhcp -chroot /var/dhc
root 87489 0.0 0.0 6152 1924 - Ss 02:03 0:00.00 /usr/local/sbin/dhcpleases6 -c /usr/local/bin
/php-cgi -f /usr/loca
root 87637 0.0 0.1 10440 2520 - Ss 02:02 0:00.02 /usr/sbin/syslogd -s -c -c -l /var/dhcpd/var/
run/log -P /var/run/s
```

**#58 - 01/23/2017 11:11 AM - Daryl Morse**

Barring unforeseen changes, you should not expect the behaviour to change between now and the release version. I've been using various versions of this fix for months on Telus without problem, so it should work for you also. What modem are you using? Is the port in bridge mode? Apparently the V2000H has some issues with bridge mode. Can you put Wireshark on the wan interface? You should see dhcpv6 solicit, advertise, request, reply and icmpv6 rs and ra. It might be better for this conversation to take place on the pfSense forum.

**#59 - 01/25/2017 11:08 AM - Martin Wasley**

And the fixes are for dhcp6c the WAN client, not dhcpd the LAN server. Different things.

First, enable dhcp6c debug and reboot and see if your dhcp6c client obtains an address and/or prefix, that should show in the dhcp logs, then in the system logs you should see something like this:

```
Jan 25 15:38:22 gateway rtsold: Received RA specifying route fe80::*:feb1:2180 for interface wan(igb0)
Jan 25 15:38:22 gateway php-fpm50330: /rc.newwanipv6: rc.newwanipv6: Info: starting on igb0.
Jan 25 15:38:22 gateway php-fpm50330: /rc.newwanipv6: rc.newwanipv6: on (IP address: fe80::*:f6:1905%igb0) (interface: wan) (real interface: igb0).
Jan 25 15:38:23 gateway dhcpleases: /etc/hosts changed size from original!
Jan 25 15:38:23 gateway dhcpleases: Could not deliver signal HUP to process because its pidfile (/var/run/unbound.pid) does not exist, No such process.
Jan 25 15:38:24 gateway dhcpleases: kqueue error: unkown
Jan 25 15:38:25 gateway php-fpm50330: /rc.newwanipv6: ROUTING: setting default route to 87.*.*.225
Jan 25 15:38:25 gateway php-fpm50330: /rc.newwanipv6: ROUTING: setting IPv6 default route to fe80::*:feb1:2180%igb0
Jan 25 15:38:25 gateway php-fpm50330: /rc.newwanipv6: Removing static route for monitor 212.58.244.22 and adding a new route through 87.*.*.*
Jan 25 15:38:25 gateway php-fpm50330: /rc.newwanipv6: Removing static route for monitor fe80::*:feb1:2180 and adding a new route through fe80::*:feb1:2180%igb0
Jan 25 15:38:25 gateway check_reload_status: Reloading filter
```

Do you?

**#60 - 01/27/2017 06:48 PM - J L**

Martin.

I got a packet capture from the WAN upon bootup (all looks good there), a copy of the system and DHCP logs.

What's the best way to contact you so I can send you a URL to a .tar on my server containing the files (as I would like to send it to you privately)?

Thanks

**#61 - 01/28/2017 08:09 AM - Martin Wasley**

Just upload them to a dropbox and send me a link.

**#62 - 01/28/2017 10:30 PM - J L**

Martin Wasley wrote:

Just upload them to a dropbox and send me a link.

Sure thing. Can I email you a private link?

**#63 - 01/29/2017 08:18 AM - Martin Wasley**

Just post the dropbox link here..

**#64 - 01/29/2017 03:45 PM - J L**

Martin Wasley wrote:

Just post the dropbox link here..

Fair enough. I looked through the pcap and there's nothing really sensitive in there. I've already changed IP's anyway as it changes upon changing WAN MAC (IPv4) and DUID (and possibly LAN MAC) for IPv6

<https://www.dropbox.com/sh/qti4eps0r8f04j9/AADeZKTIQHngScPrErdBVL4Ya?dl=0>

The logs weren't taken from the same time as the WAN pcap which I do apologize for.

Part of the issue I think is that the DHCPv6 server doesn't properly bind on the IPv6 interface upon bootup, ie

```
Jan 28 20:26:52 router dhcpd: Listening on BPF/em3/00:24:81:7e:38:aa/192.168.1.0/24
Jan 28 20:26:52 router dhcpd: Sending on   BPF/em3/00:24:81:7e:38:aa/192.168.1.0/24
Jan 28 20:26:52 router dhcpd: Sending on   Socket/fallback/fallback-net
```

vs

```
Jan 28 20:28:14 router dhcpd: Listening on Socket/5/em3/2001:569:7041:bb00::/64
Jan 28 20:28:14 router dhcpd: Sending on   Socket/5/em3/2001:569:7041:bb00::/64
```

after disabling and re-enabling the DHCPv6 server.

**#65 - 01/30/2017 03:06 AM - Martin Wasley**

OK, I can see dhcp6c is doing its job and launching RTSOLD, which is launching rc.newwanipv6. As no-one else has reported this problem and I cannot re-create it, then I need to know more about your configuration. What hardware are you running on, number of ports etc.

One of my testers is also running on Telus and has not reported this problem. Can you upload your config.xml to dropbox too so I can load it and see if there is anything there you may have set that could cause this issue. You can delete the existing files on dropbox now.

**#66 - 01/31/2017 11:01 PM - J L**

Martin Wasley wrote:

OK, I can see dhcp6c is doing its job and launching RTSOLD, which is launching rc.newwanipv6. As no-one else has reported this problem and I cannot re-create it, then I need to know more about your configuration. What hardware are you running on, number of ports etc.

One of my testers is also running on Telus and has not reported this problem. Can you upload your config.xml to dropbox too so I can load it and see if there is anything there you may have set that could cause this issue. You can delete the existing files on dropbox now.

Alright. I doubt hardware matters too much, but I'll list some facts.

Motherboard: Intel DH55HC

WAN NIC: Using integrated Intel NIC on board.

LAN NIC: 4-port HP branded Intel chipset NIC with LAN port going to Mikrotik 24-port switch (no funky VLAN/filtering stuff going on there)

WAN gateway: Actiontec V2000H with config patch to avoid the device hogging a IPv6 prefix to itself (bug, and I can confirm it's fixed)

I'd be willing to send you my config.xml privately as it contains some personal information, passwords of installed packages, etc. Suppose I could manually redact it though.

Maybe I should do a realtime video during a pfSense restart showing the issue in action after a reboot, etc.

**#67 - 02/01/2017 04:24 PM - Martin Wasley**

I'll pm you with my email address.

Edit.. you're not showing yours either.. :)

**#68 - 02/01/2017 04:27 PM - Martin Wasley**

Do you have Daryl's pm?

**#69 - 02/01/2017 04:31 PM - J L**

Martin Wasley wrote:

I'll pm you with my email address.

Edit.. you're not showing yours either.. :)

jtl at teamclassified dot ca

**#70 - 02/01/2017 06:46 PM - Daryl Morse**

J L wrote:

Alright. I doubt hardware matters too much, but I'll list some facts.

WAN gateway: Actiontec V2000H with config patch to avoid the device hogging a IPv6 prefix to itself (bug, and I can confirm it's fixed)

Is there any way you can get a different modem? I didn't see your pcap, but unless Martin is convinced it's okay, the modem is a question, especially since you're saying it has issues with the prefix. The Telus edge routers are fussy. To ensure that the WAN side is configuring properly, you should see a sequence of dhcp6c (x4) and icmpv6 (x2) messages. If the WAN side doesn't configure, the LAN won't work either.

**#71 - 02/01/2017 07:35 PM - J L**

Daryl Morse wrote:

J L wrote:

Alright. I doubt hardware matters too much, but I'll list some facts.

WAN gateway: Actiontec V2000H with config patch to avoid the device hogging a IPv6 prefix to itself (bug, and I can confirm it's fixed)

Is there any way you can get a different modem? I didn't see your pcap, but unless Martin is convinced it's okay, the modem is a question, especially since you're saying it has issues with the prefix. The Telus edge routers are fussy. To ensure that the WAN side is configuring properly, you should see a sequence of dhcp6c (x4) and icmpv6 (x2) messages. If the WAN side doesn't configure, the LAN won't work either.

I believe there's two things of interest going on here.

1. The Actiontec possibly being the "link-local" next-hop for IPv6, which could cause issues and explains my 0ms gateway time on gateway monitoring, although since IPv6 works once I poke the dhcp6c server. I doubt it's causing the current issue but could be wrong. I looked through the pcap of my WAN and see the dhcpv6 and icmpv6 packets as mentioned.
2. The dhcp server possibly not binding to IPv6 at startup

I made a video showing a restart of pfSense, and tailing the DHCP logfile in real time while disabling and re-enabling the dhcp6c server.

I might look into getting a new modem if it comes to that, but also I might getting FTTH in a few months.

<https://www.youtube.com/watch?v=yrFD-pYbAwg>

J L wrote:

Daryl Morse wrote:

J L wrote:

Alright. I doubt hardware matters too much, but I'll list some facts.

WAN gateway: Actiontec V2000H with config patch to avoid the device hogging a IPv6 prefix to itself (bug, and I can confirm it's fixed)

Is there any way you can get a different modem? I didn't see your pcap, but unless Martin is convinced it's okay, the modem is a question, especially since you're saying it has issues with the prefix. The Telus edge routers are fussy. To ensure that the WAN side is configuring properly, you should see a sequence of dhcp6c (x4) and icmpv6 (x2) messages. If the WAN side doesn't configure, the LAN won't work either.

I believe there's two things of interest going on here.

1. The Actiontec possibly being the "link-local" next-hop for IPv6, which could cause issues and explains my 0ms gateway time on gateway monitoring, although since IPv6 works once I poke the dhcp6c server. I doubt it's causing the current issue but could be wrong. I looked through the pcap of my WAN and see the dhcpv6 and icmpv6 packets as mentioned.
2. The dhcp server possibly not binding to IPv6 at startup

I made a video showing a restart of pfSense, and tailing the DHCP logfile in real time while disabling and re-enabling the dhcp6c server.

I might look into getting a new modem if it comes to that, but also I might getting FTTH in a few months.

<https://www.youtube.com/watch?v=yrED-pYbAwg>

There is definitely something going on with the short hop for the dhcp6 gateway. Both of my gateways are around 6 ms. My modem is completely transparent. I don't think you can say the same thing for yours. Even if you're going to get FTTH, you will still need an actiontec. It will just connect to the FTTH box using ethernet. There's no way to guarantee it will make a difference, however. Also, I'm not using any apple computers, so no way to verify if they work properly on my network.

**#73 - 02/02/2017 04:52 AM - Martin Wasley**

It's not his modem, he's getting a prefix, all is well there. Just looked at the video. Let me think on it...

Going to send you a patch and client. It's not dhcp6c causing the issue but the patch needs the client I am sending in order to work. Copy the client into /usr/local/sbin and apply the patch, clear your system and dhcp logs and then reboot. Once the system is back up wait about 60 seconds then save the logs and send them to me. You'll have my email address.

Also I noticed you have two LAN ports running, let's make sure that is not a cause of the issue. After you have tested with the patch and client if it's still the same can you disable the 2nd lan port, reboot and see what happens then.

**#74 - 02/10/2017 10:27 AM - Jim Pingle**

- Target version changed from 2.4.0 to 2.3.3

**#75 - 02/10/2017 11:33 AM - Renato Botelho**

- Status changed from Feedback to New

- Assignee changed from Jim Thompson to Renato Botelho

- Target version changed from 2.3.3 to 2.4.0

Not finished yet

**#76 - 02/28/2017 12:12 PM - Daryl Morse**

Renato Botelho wrote:

Not finished yet

This issue has been fixed in both 2.3.3 release and 2.4.0 beta.

**#77 - 02/28/2017 12:15 PM - Luke Hamburg**

Daryl, are you saying the marjohn56 patches are no longer needed?

**#78 - 02/28/2017 01:02 PM - Daryl Morse**

Definitely no, but the original problem for which this issue was raised has been fixed by a PR that was merged previously. The new patches are for other issues. (Sorry, not sure what numbers.)

**#79 - 02/28/2017 05:26 PM - Daryl Morse**

The new patches address issues 7145 and 7185. The other issue I was thinking of is 6944 (DHCP no release), but it has also been fixed.

**#80 - 03/16/2017 11:43 AM - Daryl Morse**

This issue can be closed. It was fixed in 2.3.3 and 2.4.

**#81 - 03/16/2017 11:45 AM - Jim Pingle**

- Status changed from New to Feedback

**#82 - 06/19/2017 06:13 AM - Martin Wasley**

This one should be closed Jim, it's been rock solid for months now.

Fixed - Resolved.

**#83 - 06/19/2017 07:16 AM - Jim Pingle**

- Status changed from Feedback to Resolved

**Files**

---

interfaces.inc.patch	849 Bytes	03/16/2016	Richard Patterson
redmine5993.patch	5.47 KB	03/18/2016	Chris Buechler
interfaces.inc.patch	1.17 KB	03/25/2016	Martin Wasley
interfaces.inc.patch	2.01 KB	04/06/2016	Martin Wasley
interfaces.php.patch	404 Bytes	04/06/2016	Martin Wasley