

pfSense - Bug #6099

igmpproxy does not recognize upstream interface

04/09/2016 04:49 PM - Victor Toni

Status:	Resolved	Start date:	04/09/2016
Priority:	Normal	Due date:	
Assignee:	Luiz Souza	% Done:	100%
Category:	IGMP Proxy	Estimated time:	0.00 hour
Target version:	2.4.0		
Affected Version:	2.3	Affected Architecture:	All

Description

igmpproxy does not start anymore with a working config from 2.2.6.

The reason seems to be that igmpproxy does not recognize the interface.

My WAN (bge0) interface is configured for VLAN 7 (PPPOE) and VLAN 8 (DHCP).

igmpproxy is configured to use bge0_vlan8 as upstream. bge0_vlan8 is not recognized, only bge0 (which shouldn't).

The main network is 10.0/16.

In the subnet 10.0.15/24 there might be some clients using VLC to request multicast from igmpproxy.

The dedicated network 192.168.11/24 is used for the set top boxes (IPTV).

The network 192.168.1/24 is used for accessing the modem.

The config was used to create the attached log by calling igmpproxy manually via
igmpproxy -d -vvvv /tmp/igmpproxy.conf

This part seems most relevant:

buildIfVc: Interface igb0 Addr: 10.0.0.1, Flags: 0xffff8943, Network: 10.0/16

buildIfVc: Interface igb2 Addr: 192.168.11.1, Flags: 0xffff8843, Network: 192.168.11/24

buildIfVc: Interface bge0 Addr: 192.168.1.11, Flags: 0xffff8943, Network: 192.168.1/24

buildIfVc: Interface lo0 Addr: 127.0.0.1, Flags: 0xffff8049, Network: 127/8

Found config for igb0

Found config for igb2

Found config for bge0

adding VIF, lx 0 FI 0x0 IP 0x0100000a igb0, Threshold: 1, Ratelimit: 0

Network for [igb0] : 10.0/16

Network for [igb0] : 10.0.15/24

adding VIF, lx 1 FI 0x0 IP 0x010ba8c0 igb2, Threshold: 1, Ratelimit: 0

Network for [igb2] : 192.168.11/24

Network for [igb2] : 192.168.11/24

adding VIF, lx 2 FI 0x0 IP 0x0b01a8c0 bge0, Threshold: 1, Ratelimit: 0

Network for [bge0] : 192.168.1/24

There must be at least 2 Vif's where one is upstream.

bge0_vlan8 does not appear.

History

#1 - 04/09/2016 05:01 PM - Chris Buechler

- Status changed from New to Confirmed

- Target version set to 2.3.1

Thanks, I had your forum thread up still to check and get a bug ticket opened. I believe it's the VLAN misinterpretation in igmpproxy.

#2 - 04/10/2016 04:56 AM - Victor Toni

There is a new home for igmpproxy at GitHub: <https://github.com/pali/igmpproxy/>

The "next" branch contains a lot of patches which were previously scattered in different places.

I believe these patches are not in the upstream package for FreeBSD.

Maybe the GitHub version would be a good starting point for further tests.

I wondering if there are also some internal changes to how fbsd handles multicast since I tried the igmpproxy from the 2.2.6 image on the 2.3 installation and it didn't work.

Unfortunately I have no logs of this test (and reverted back to 2.2.6).

#3 - 04/13/2016 02:12 PM - Pr0 vieH

I hope this is resolved quickly, it affects all "Deutsche Telekom Entertain" users they use pfsense.

#4 - 04/13/2016 04:05 PM - Stefan Heck

Yea, very annoying bug.

#5 - 04/14/2016 05:20 AM - Matthias Lange

Hi, yes indeed. This tool is needed by many IPTV users. So please fix it with a higher priority.

#6 - 04/14/2016 11:54 AM - Victor Toni

For me it's not important if the multicast part is done by igmpproxy or something else.

Actually thinking about an alternative might be good idea because igmpproxy has some shortcomings:

- supports only IPv4 (upstream as downstream)
- as far I understand it handles IGMPv3 only in parts

An alternative might be <https://github.com/mcproxy/mcproxy>

It seems that it was designed with Linux in mind.

But looking at the OS specific code parts of igmpproxy this might not that much of a difference...

(Although I'm not an expert for network code by any means and IANANE, I am not a network engineer ;-))

#7 - 04/14/2016 12:00 PM - Victor Toni

If mcproxy is not an alternative, here is the new home of igmpproxy:

<https://github.com/pali/igmpproxy>

There is also a pending pull request:

<https://github.com/pali/igmpproxy/pull/2>

which includes the interface related parts of the pfSense patch from [#5783](#):

<https://github.com/pfsense/FreeBSD-ports/commit/28cdacad47051a02e7866093634de3326242d21e>

#8 - 04/16/2016 09:02 PM - Jim Thompson

- Assignee set to Renato Botelho

#9 - 04/26/2016 11:48 PM - Chris Buechler

- Target version changed from 2.3.1 to 2.3.2

#10 - 04/27/2016 01:43 PM - Johannes Wanink

- Why is the target version now 2.3.2? :(Is there a workaround?
- Another stupid question: Is "Affected Architecture" set correct? Should it changed to "all"?

#11 - 04/27/2016 01:45 PM - Victor Toni

Tried to compile the version from the "next" branch: <https://github.com/pali/igmpproxy>

Compilation failed due to a difference between Linux and FreeBSD regarding the call of setpgrp (in igmpproxy.c:137). The Linux version expects no arguments whereas FreeBSD needs (pid_t pid, pid_t pgrp).

The solution might be something like

```
int setpgrp(void)
{
    return setpgid(0, 0);
}
```

The question is where this definition should go to since it is FreeBSD specific.

Edit:

The question has been solved. After a bit of fiddling I could create a version which compiles for FreeBSD and Linux:

<https://github.com/pali/igmpproxy/pull/3>

#12 - 04/27/2016 03:36 PM - Victor Toni

Johannes Wanink wrote:

- Another stupid question: Is "Affected Architecture" set correct? Should it changed to "all"?

I have only i386 available, so it's the only platform I could test. Maybe it should be changed to all.

#13 - 04/27/2016 04:42 PM - Chris Buechler

- Affected Architecture changed from i386 to All

Johannes Wanink wrote:

- Why is the target version now 2.3.2? :(Is there a workaround?

Because there is no fix available and we're releasing 2.3.1 soon. No known workaround.

it does affect all architectures

#14 - 04/28/2016 09:15 AM - Stefan Heck

- File *router.skeba.de - Status_Dashboard_2016-04-28_16-14-15.png* added

It works even with pfSense 2.3 after I joined Receiver and other Network Hardware to the same LAN-Subnet.
So pfSense 2.3 works on my APU 1d4 with Entertain for 3 days now

#15 - 04/28/2016 09:23 AM - Matthias Lange

Hello Stefan,
could you please go more in detail on this? Maybe some screenshots. I don't know what you mean with "receiver".

#16 - 04/28/2016 10:54 AM - Stefan Heck

- File *Dok2.pdf* added

Basically the thread starter is Entertain-Customer of the Deutsche Telekom, so I am – IPTV Entertain is using Media Receiver to playback Multicast streams on TV. For this issue this is not of any interest – sorry for using the term "Receiver" – use it as a synonym for the soft/hardware you are using to Playback Multicaststreams in combination with the igmpproxy.

My Primary intention was to separate LAN and IPTV by using different interfaces/subnet'S – WAN, LAN, IPTV. This failed!

Now I have only WAN and LAN interface active and physically attached my playback hardware on the switch behind my LAN-Interface.

I receive "Internet" from my ISP via DSL using VLAN7 and "IPTV" using VLAN8

in the attached Document you will find some screenshots - but there is nothing special in this configuration - it simply works as expected - dont know why.

```
[2.3-RELEASE][root@router]:root: cat /tmp/igmpproxy.conf  
##-----
```

```
1. Enable Quickleave mode (Sends Leave instantly)
```

```
##-----  
quickleave  
phyint re0_vlan8 upstream ratelimit 0 threshold 1  
altnet 193.158.0.0/15  
altnet 224.0.0.0/4
```

```
phyint re1 downstream ratelimit 0 threshold 1  
altnet 192.168.1.0/24
```

```
phyint pppoe0 disabled
```

#17 - 04/28/2016 11:11 AM - Matthias Lange

Hi Stefan,

this is exactly my config since month for IPTV with DeTelekom and Entertain.
So maybe the config from Victor was not this way? What was your config before?

Here is my conf file

```
##-----  
  
1. Enable Quickleave mode (Sends Leave instantly)  
##-----  
quickleave  
phyint em1_vlan8 upstream ratelimit 0 threshold 1  
altnet 239.35.0.0/16  
altnet 217.0.119.0/24  
altnet 193.158.35.0/24  
altnet 10.0.0.0/8  
  
phyint em0 downstream ratelimit 0 threshold 1  
altnet 192.168.20.0/24  
  
phyint pppoe0 disabled
```

#18 - 04/28/2016 11:25 AM - Stefan Heck

Unfortunately, I cannot provide my previous config, I have started with pfSense two weeks ago und struggled directly into the migration process from 2.2.6 to 2.3.

I only know that it had worked with 2.2.6 and stopped with 2.3

After a complete reinstall, it is now working. Maybe some wrong parameter left after the update prozess?

#19 - 04/28/2016 11:28 AM - Matthias Lange

ok, can we talk offline? here is my email: matthias@lange-sd.de

#20 - 04/28/2016 11:33 AM - Stefan Heck

@Matthias Lange
you have mail

#21 - 04/28/2016 01:27 PM - Victor Toni

Stefan Heck wrote:

After a complete reinstall, it is now working. Maybe some wrong parameter left after the update prozess?

This might be a valid point, I've done only an upgrade from 2.2.6 so it could be related to the upgrade.

#22 - 04/28/2016 01:32 PM - Victor Toni

Stefan Heck wrote:

```
phyint re0_vlan8 upstream ratelimit 0 threshold 1
...
altnet 224.0.0.0/4
```

Matthias Lange wrote:

```
phyint em1_vlan8 upstream ratelimit 0 threshold 1
...
altnet 10.0.0.0/8
```

Maybe someone could explain to me why these entries are on the external/upstream interface?

In my understanding only the *real* addresses for the stream sources need to be mentioned in the upstream section.

#23 - 04/29/2016 02:29 PM - Ulysse FONTAINE

I have the same problem, my ISP use VLAN 838, 840, 839 for IPTV. I have made a bridge with thoes vlans, i have an ip on the bridge. But the IcmpProxy service gives me:

```
Apr 29 21:16:00 igmpproxy 1158 There must be at least 2 Vif's where one is upstream.
Apr 29 21:16:00 igmpproxy 1158 adding VIF, Ix 3 F1 0x0 IP 0xfe0b010a lagg0_vlan11, Threshold: 1, Ratelimit: 0
Apr 29 21:16:00 igmpproxy 1158 adding VIF, Ix 2 F1 0x0 IP 0xfe01010a lagg0, Threshold: 1, Ratelimit: 0
Apr 29 21:16:00 igmpproxy 1158 adding VIF, Ix 1 F1 0x0 IP 0xfe00a8c0 em2, Threshold: 1, Ratelimit: 0
Apr 29 21:16:00 igmpproxy 1158 adding VIF, Ix 0 F1 0x0 IP 0x320aa8c0 re0, Threshold: 1, Ratelimit: 0
Apr 29 21:16:00 php-fpm 143 /status_services.php: Started IGMP proxy service.
```

Here are my VLANS:

```
lagg0 (opt4) 11 MailNetwork
lagg0 (opt4) 10 WebNetwork
lagg0 (opt4) 2 ServiceNetwork
lagg0 (opt4) 100 Media
re0 (opt6) 838 TV DHCP
re0 (opt6) 840 TV STREAM
re0 (opt6) 835 WAN
```

And my igmpProxy config:

```
TV      upstream      193.0.0.0/8, 81.0.0.0/8, 172.0.0.0/8, 80.0.0.0/8      TV_ORANGE_UP
LAN     downstream      192.168.0.0/24      TV_ORANGE_DOWN
```

I'm on a fresh config of pfsense 2.3. I upgraded my router but I have decided to make a fresh install when it came out.

If you want more information you can mail me at: ulyse@nored.ovh or respond to this issue.

By the way, i heard that this probleme can come from a loop or a counting loop on the igmpProxy program. Because some peoples can run igmpProxy with the first 3 vlans.

Here is the source: <https://forum.pfsense.org/index.php?topic=106153.0>

#24 - 04/30/2016 01:13 AM - Down Loadski

Stefan Heck wrote:

Basically the thread starter is Entertain-Customer of the Deutsche Telekom, so I am – IPTV Entertain is using Media Receiver to playback Multicast streams on TV. For this issue this is not of any interest – sorry for using the term “Receiver” – use it as a synonym for the soft/hardware you are using to Playback Multicaststreams in combination with the igmpProxy.

My Primary intention was to separate LAN and IPTV by using different interfaces/subnet'S – WAN, LAN, IPTV. This failed!

Now I have only WAN and LAN interface active and physically attached my playback hardware on the switch behind my LAN-Interface.

I receive “Internet” from my ISP via DSL using VLAN7 and “IPTV” using VLAN8

Thanks for the info.

Are you saying that the LAN side needs to be physical interface, and cannot be a vlan ? The WAN side is still working with vlans with you, that is great as that would be hard to overcome.

So a workaround would than be to make a LAN connection with the settopbox in it, and have an additional intereface from the pfsense box with all other internal vlans on it if needed.

It would be nice to split the IPTV from normal LAN with all the multicast traffic, but if it needs to be in there, so be it. iGMP snooping on all the interfaces on the LAN will help.

#25 - 04/30/2016 05:21 AM - Kay Ringmann

I have a APU1D4.

re0 = WAN

re1,re2 = LAGG0

LAGG0 = VLAN1,VLAN10,VLAN20,VLAN30,VLAN40,VLAN50,VLAN60

So the interfaces are represented as LAGG0_VLAN10, LAGG0_VLAN20 ...

The LAGG0 are connected to a L2 switch (HP-1820) where i split the VLAN's to different ports. The IGMPProxy has to work between LAGG0_VLAN50 and LAGG0_VLAN60.

```
##-----
## Enable Quickleave mode (Sends Leave instantly)
##-----
quickleave
phyint lagg0_vlan50 upstream ratelimit 0 threshold 1
altnet 192.168.50.1/24

phyint lagg0_vlan60 downstream ratelimit 0 threshold 1
altnet 192.168.60.1/24

phyint re0 disabled
phyint lagg0_vlan1 disabled
phyint lagg0_vlan10 disabled
phyint lagg0_vlan20 disabled
phyint lagg0_vlan30 disabled
phyint lagg0_vlan40 disabled

# igmpproxy -d -v /tmp/igmpproxy.conf
adding VIF, Ix 0 Fl 0x0 IP 0xd7e1111f re0, Threshold: 1, Ratelimit: 0
adding VIF, Ix 1 Fl 0x0 IP 0x010110ac lagg0_vlan1, Threshold: 1, Ratelimit: 0
adding VIF, Ix 2 Fl 0x0 IP 0x010aa8c0 lagg0_vlan10, Threshold: 1, Ratelimit: 0
adding VIF, Ix 3 Fl 0x0 IP 0x0114a8c0 lagg0_vlan20, Threshold: 1, Ratelimit: 0
There must be at least 2 Vif's where one is upstream.
#
```

As you see, not all (more the up-counted 4) VLAN's are recognized. For testing i changed the config:

```
##-----
## Enable Quickleave mode (Sends Leave instantly)
##-----
quickleave
phyint lagg0_vlan10 upstream ratelimit 0 threshold 1
altnet 192.168.10.1/24

phyint lagg0_vlan20 downstream ratelimit 0 threshold 1
altnet 192.168.20.1/24

phyint re0 disabled
phyint lagg0_vlan1 disabled
phyint lagg0_vlan30 disabled
phyint lagg0_vlan40 disabled
phyint lagg0_vlan50 disabled
phyint lagg0_vlan60 disabled
```

and the igmpproxy comes up:

```
adding VIF, Ix 0 Fl 0x0 IP 0xd7e1111f re0, Threshold: 1, Ratelimit: 0
adding VIF, Ix 1 Fl 0x0 IP 0x010110ac lagg0_vlan1, Threshold: 1, Ratelimit: 0
adding VIF, Ix 2 Fl 0x0 IP 0x010aa8c0 lagg0_vlan10, Threshold: 1, Ratelimit: 0
adding VIF, Ix 3 Fl 0x0 IP 0x0114a8c0 lagg0_vlan20, Threshold: 1, Ratelimit: 0
joinMcGroup: 224.0.0.2 on lagg0_vlan20
```

...but where are the IF's lagg0_vlan30 till lagg0_vlan60 ???

The igmpproxy worked like a charme till 2.2.6.

While reading many posts related to this problem, i'dont think it's a problem of naming the interfaces, rather then the number of them.

Hope this helps for solving the problem.

#26 - 04/30/2016 05:43 AM - Victor Toni

Kay Ringmann wrote:

While reading many posts related to this problem, i'dont think it's a problem of naming the interfaces, rather then the number of them.

Good point. Checking the source there seems to be enough headroom for some more interfaces
<https://github.com/pali/igmpproxy/blob/master/src/igmpproxy.h>

Line 113:

```
#define MAX_IF 40 // max. number of interfaces recognized
```

so maybe it's something while parsing all the interfaces not the number itself?

#27 - 05/01/2016 11:31 AM - Kay Ringmann

i read this. So I played a little around. It seems that it's depending from upstream interface.

In my earlier posts i changed up- and downstream if's less than 4 of counted if's. for now i only changed the upstream down to the 4 first if's tagged with vlan.

the result:

upstream lagg0_vlan10, downstream lagg0_vlan60 => working

upstream lagg0_vlan20, downstream lagg0_vlan60 => working

upstream lagg0_vlan30, downstream lagg0_vlan60 => NOT WORKING

changed upstream to lagg0_vlan10 and downstream to ALL other vlans (20,30,40,50,60) => working

can anybody confirm that behavior?

#28 - 05/02/2016 01:26 AM - Down Loadski

Kay Ringmann wrote:

i read this. So I played a little around. It seems that it's depending from upstream interface.

In my earlier posts i changed up- and downstream if's less than 4 of counted if's. for now i only changed the upstream down to the 4 first if's tagged with vlan.

the result:

upstream lagg0_vlan10, downstream lagg0_vlan60 => working

upstream lagg0_vlan20, downstream lagg0_vlan60 => working

upstream lagg0_vlan30, downstream lagg0_vlan60 => NOT WORKING

changed upstream to lagg0_vlan10 and downstream to ALL other vlans (20,30,40,50,60) => working

can anybody confirm that behavior?

Interesting..

The error:

"There must be at least 2 Vif's where one is upstream."

Is coming from the igmpmproxy.c source only if i checked well enough.

I went through many of the config of the source on github and it seems to interpret the config file and determines out of that the downstream and upstream interfaces and stores that in an internal structure, which get referenced and than errors out when no 2 vif in total, or no upstream vif is found.

So a conclusion might be that it does not read a correct upstream vif at all in some conditions from the config file.

I will try to test it tomorrow when the wife and kids are not at home.

#29 - 05/07/2016 07:16 AM - Ulysse FONTAINE

Kay Ringmann wrote:

i read this. So I played a little around. It seems that it's depending from upstream interface.

In my earlier posts i changed up- and downstream if's less than 4 of counted if's. for now i only changed the upstream down to the 4 first if's tagged with vlan.

the result:

upstream lagg0_vlan10, downstream lagg0_vlan60 => working

upstream lagg0_vlan20, downstream lagg0_vlan60 => working

upstream lagg0_vlan30, downstream lagg0_vlan60 => NOT WORKING

changed upstream to lagg0_vlan10 and downstream to ALL other vlans (20,30,40,50,60) => working

can anybody confirm that behavior?

I confirm, if you take interfaces that are listed, igmp proxy work otherwise it doesn't

the only interfaces he could detect was:

```
buildIfVc: Interface re0 Addr: 192.168.10.50, Flags: 0xffff8943, Network: 192.168.10/24
buildIfVc: Interface em2 Addr: 192.168.0.254, Flags: 0xffff8843, Network: 192.168.0/24
buildIfVc: Interface lo0 Addr: 127.0.0.1, Flags: 0xffff8049, Network: 127/8
buildIfVc: Interface lagg0 Addr: 10.1.1.254, Flags: 0xffff8843, Network: 10.1.1/24
buildIfVc: Interface lagg0_vlan11 Addr: 10.1.11.254, Flags: 0xffff8843, Network: 10.1.11/24
```

so using re0 as upstream and em2 as downstream, igmp proxy work.

here is the config file:

```
phyint re0 upstream
    altnet 193.0.0.0/8 altnet 81.0.0.0/8 altnet 172.0.0.0/8 altnet 80.0.0.0/8

phyint em2 downstream
```

#30 - 05/07/2016 04:16 PM - Kay Ringmann

I Did a Plain install, but the problem persists.

The behavior of listed Interfaces changes unpredictable if i Change the upstream Interfaces. I can't Change the upstream Interface to an other one. IT already worked on 2.2.6.

What was changed in the new Release? I dont think i'm the only one with this Problem.

#31 - 05/07/2016 05:18 PM - Chris Buechler

Kay Ringmann wrote:

What was changed in the new Release? I dont think i'm the only one with this Problem.

Updated to latest igmp proxy version, and FreeBSD 10.3 base. This is in one or the combination of those. Much if not all of it is in igmp proxy itself.

#32 - 05/10/2016 02:52 AM - Kay Ringmann

For testing I copied the igmpproxy from version 2.2.6 to latest pfsense Version. Then I had to Link /tmp/igmpproxy.conf with /usr/local/etc/igmpproxy.conf and it started up without any problems.

#33 - 05/12/2016 04:21 PM - Lars Karow

Seems like, that ifvc.c (igmpproxy) does NOT create the full list of all interfaces in function "buildIfVc()".

Maybe there is someone who can fix the usage of "ioctl(Sock, SIOCGIFCONF, &IoCtlReq)" or at least knows how to do the correct size calculation for "IfEp = (void *)((char *)IfVc + IoCtlReq.ifc_len)". Which seems to be too small.

On github I found a patch for an old version, which removed the parts using " ioctl(Sock, SIOCGIFCONF, &IoCtlReq)" and replaced it with "getifaddrs".

Maybe this patch can be adapted to the used version of igmpproxy. Please have a look at:

<https://github.com/ironbits/pfsense-tools/blob/master/pfPorts/igmpproxy/files/patch-fbsd> starting at line 242.

#34 - 05/13/2016 04:21 PM - Lars Karow

ioctl(Sock, SIOCGIFCONF, &IoCtlReq) returns a lot more interfaces in FreeBSD 10.3 than my tests under Debian showed, for a comparable setup. For each interface up to 3 values are returned, while under Debian only interfaces with IPs are returned.

Under FreeBSD/pfSense a setup with 11 or more (VLAN-)interfaces can result in not finding all interfaces.

#35 - 05/17/2016 07:20 AM - Ulysse FONTAINE

Hi, I have some news.

Someone have made a commit on the pali/igmpproxy git: <https://github.com/pali/igmpproxy/pull/2> | <https://redmine.pfsense.org/issues/5783>
And apparently igmpproxy work fine with this fix. I didn't tested it.

But for me i used a previous version of igmpproxy on a pfSense-2.2.6 32-bit. The igmpproxy service work and start.

I used Wireshark to see the igmpproxy packets, and i noticed that IGMPv3 packets are corrupt. I don't know if it comes from the igmpproxy but it's starting at least.

To recap. I transfered igmpproxy from pfSense 2.2.6 to 2.3 and the service started.

#36 - 05/20/2016 05:33 AM - Victor Toni

Ulysse FONTAINE wrote:

Someone have made a commit on the pali/igmpproxy git: <https://github.com/pali/igmpproxy/pull/2> | <https://redmine.pfsense.org/issues/5783>

That someone was me (see [#6099#note-7](#)). :-)

I patched an up to date branch with the patch from:

<https://github.com/pfsense/FreeBSD-ports/commit/28cdacad47051a02e7866093634de3326242d21e>
(for [#5783](#))

This is different to <https://github.com/ironbits/pfsense-tools/blob/master/pfPorts/igmpproxy/files/patch-fbsd>

I can code some C but am currently lost when it comes to networking & C.

So maybe someone could take deeper look at the code.

#37 - 05/26/2016 04:56 PM - Victor Toni

Lars Karow wrote:

Seems like, that ifvc.c (igmpproxy) does NOT create the full list of all interfaces in function "buildIfVc()".

Maybe there is someone who can fix the usage of "ioctl(Sock, SIOCGIFCONF, &IoCtlReq)" or at least knows how to do the correct size calculation for "IfEp = (void *)((char *)IfVc + IoCtlReq.ifc_len)". Which seems to be too small.

Thanks for the hint. Following this path I found the reason to be in <https://github.com/pali/igmpproxy/blob/next/src/ifvc.c#L96>

```
if ( IfPt->ifr_addr.sa_family != AF_INET ) {  
    if (Dp == IfDescEp) {  
        IfDescEp++;  
    }  
    Dp->InAdr.s_addr = 0; /* mark as non-IP interface */  
    continue;  
}
```

In my case the VLAN device is recognized as non-IP and thus ignored later in the process...

The reason is that the sa_family for this interface are only AF_LINK and AF_INET6 and no AF_INET.

#38 - 05/27/2016 05:48 AM - Philipp Resch

This might be of interest, if not, please remove.

I am already on the "new" German Telekom platform, where VLAN8 no longer exists. Therefore I need to create an upstream interface on my pppoe interface.

```
igmpproxy.conf:  
@  
cat /tmp/igmpproxy.conf
```

##-----

1. Enable Quickleave mode (Sends Leave instantly)

```
##-----
quickleave
phyint pppoe1 upstream ratelimit 0 threshold 1
altnet 193.158.0.0/15
altnet 224.0.0.0/4
```

```
phyint em1 downstream ratelimit 0 threshold 1
altnet 192.168.1.0/24
```

```
phyint pppoe0 disabled
phyint em0 disabled
phyint gif0 disabled
phyint em2 disabled
@
```

Although igmpproxy seems not to be able to detect the pppoe interface:

```
igmpproxy -d -vvv /tmp/igmpproxy.conf
adding VIF, Ix 0 FI 0x0 IP 0x0202a8c0 em0, Threshold: 1, Ratelimit: 0
adding VIF, Ix 1 FI 0x0 IP 0x0101a8c0 em1, Threshold: 1, Ratelimit: 0
adding VIF, Ix 2 FI 0x0 IP 0xfd0110ac em2, Threshold: 1, Ratelimit: 0
There must be at least 2 Vif's where one is upstream.
```

Maybe this issue is related, as it seems as if pppoe interfaces are also not detected by the buildIfVc() function.
I fiddled around with it a bit, by adding AF_LINK to be a valid interface, but that breaks quite a lot of other stuff...

#39 - 05/27/2016 05:54 PM - Victor Toni

Lars Karow wrote:

Seems like, that ifvc.c (igmpproxy) does NOT create the full list of all interfaces in function "buildIfVc()".

Maybe there is someone who can fix the usage of "ioctl(Sock, SIOCGIFCONF, &IoCtlReq)" or at least knows how to do the correct size calculation for "IfEp = (void *)((char *)IfVc + IoCtlReq.ifc_len)". Which seems to be too small.

On github I found a patch for an old version, which removed the parts using " ioctl(Sock, SIOCGIFCONF, &IoCtlReq)" and replaced it with "getifaddrs".

Maybe this patch can be adapted to the used version of igmpproxy. Please have a look at:

<https://github.com/ironbits/pfsense-tools/blob/master/pfPorts/igmpproxy/files/patch-fbsd> starting at line 242.

Trying to make `icoctl()` work as in <http://stackoverflow.com/a/4139564> (including comment) was not successful, maybe due to <http://stackoverflow.com/questions/7629164/enumerate-all-network-interfaces-with-ips-on-freebsd> (first comment).

So I went for the `getifaddrs()` approach and integrated the patch into my branch. It seems to detect now all the interfaces.

<https://github.com/ViToni/igmpproxy/commit/79f43a55095c466cac475ef959e0e4447c0e285b>

Testers welcome ;-)

There might be other issues, like dynamic interface detection or whitelists broken, but it could be a start.

#40 - 05/28/2016 09:03 AM - Greg Myran

Victor Toni wrote:

So I went for the `getifaddrs()` approach and integrated the patch into my branch. It seems to detect now all the interfaces.

<https://github.com/ViToni/igmpproxy/commit/79f43a55095c466cac475ef959e0e4447c0e285b>

Testers welcome ;-)

There might be other issues, like dynamic interface detection or whitelists broken, but it could be a start.

With this patched version I now see all interfaces being correctly detected and properly assigned to upstream/downstream. The IGMP joins are sent on the upstream interface and the multicast flow starts but is not forwarded to the downstream interface and times out after 2 seconds.

```
RECV V2 member report from 10.0.3.10 to 239.0.0.35
Should insert group 239.0.0.35 (from: 10.0.3.10) to route table. Vif ix : 0
No existing route for 239.0.0.35. Create new.
No routes in table. Insert at beginning.
Inserted route table entry for 239.0.0.35 on VIF #0
Joining group 239.0.0.35 upstream on IF address 10.101.0.2
joinMcGroup: 239.0.0.35 on gre1
```

Current routing table (Insert Route):

```
-----
#0: Dst: 239.0.0.35, Age:2, St: I, OutVifs: 0x00000001
-----
```

```
RECV V2 member report from 10.101.0.2 to 239.0.0.35
The IGMP message was from myself. Ignoring.
Route activate request from 192.168.0.35 to 239.0.0.35
Vif bits : 0x00000001
Setting TTL for Vif 0 to 0
Setting TTL for Vif 0 to 1
Setting TTL for Vif 0 to 0
```

Adding MFC: 192.168.0.35 -> 239.0.0.35, InpVif: 17

Current routing table (Activate Route):

#0: Src0: 192.168.0.35, Dst: 239.0.0.35, Age:2, St: A, OutVifs: 0x00000001

About to call timeout 10 (#0)
SENT Membership query from 10.0.3.1 to 224.0.0.1
Sent membership query from 10.0.3.1 to 224.0.0.1. Delay: 10
Created timeout 13 (#0) - delay 10 secs
(Id:13, Time:10)
Created timeout 14 (#1) - delay 115 secs
(Id:13, Time:10)
(Id:14, Time:115)
About to call timeout 13 (#0)
Aging routes in table.

Current routing table (Age active routes):

#0: Src0: 192.168.0.35, Dst: 239.0.0.35, Age:1, St: A, OutVifs: 0x00000001

About to call timeout 14 (#0)
SENT Membership query from 10.0.3.1 to 224.0.0.1
Sent membership query from 10.0.3.1 to 224.0.0.1. Delay: 10
Created timeout 15 (#0) - delay 10 secs
(Id:15, Time:10)
Created timeout 16 (#1) - delay 115 secs
(Id:15, Time:10)
(Id:16, Time:115)
About to call timeout 15 (#0)
Aging routes in table.
Removing group 239.0.0.35. Died of old age.
Removed route entry for 239.0.0.35 from table.
Vif bits : 0x00000001
Setting TTL for Vif 0 to 0
Setting TTL for Vif 0 to 1
Setting TTL for Vif 0 to 0
Removing MFC: 192.168.0.35 -> 239.0.0.35, InpVif: 17
Leaving group 239.0.0.35 upstream on IF address 10.101.0.2
leaveMcGroup: 239.0.0.35 on gre1

Current routing table (Remove route):

No routes in table...

#41 - 05/30/2016 03:41 AM - Philipp Resch

Hi, your patched version does now also find the pppoe interfaces, although I was not yet able to get a picture.

Here's a log:

```
./igmpproxy -d -vvv /tmp/igmpproxy.conf
adding VIF, Ix 0 FI 0x0 IP 0x0101a8c0 em1, Threshold: 1, Ratelimit: 0
adding VIF, Ix 1 FI 0x0 IP 0xcb8ee95d pppoe1, Threshold: 1, Ratelimit: 0
joinMcGroup: 224.0.0.2 on em1
joinMcGroup: 224.0.0.22 on em1
RECV V3 member report from 192.168.1.1 to 224.0.0.22
The IGMP message was from myself. Ignoring.
The IGMP message was from myself. Ignoring.
RECV V3 member report from 192.168.1.1 to 224.0.0.22
The IGMP message was from myself. Ignoring.
The IGMP message was from myself. Ignoring.
RECV V3 member report from 192.168.1.222 to 224.0.0.22
Inserted route table entry for 232.0.20.35 on VIF #0
joinMcGroup: 232.0.20.35 on pppoe1
Adding MFC: 87.141.215.251 -> 232.0.20.35, InpVIf: 15
RECV V3 member report from 192.168.1.222 to 224.0.0.22
Updated route entry for 232.0.20.35 on VIF #0
Adding MFC: 87.141.215.251 -> 232.0.20.35, InpVIf: 15
RECV V3 member report from 192.168.1.222 to 224.0.0.22
Updated route entry for 232.0.20.35 on VIF #0
Adding MFC: 87.141.215.251 -> 232.0.20.35, InpVIf: 15
RECV V3 member report from 192.168.1.222 to 224.0.0.22
RECV V3 member report from 192.168.1.222 to 224.0.0.22
RECV V3 member report from 192.168.1.222 to 224.0.0.22
```

#42 - 05/30/2016 04:40 AM - Victor Toni

Philipp Resch wrote:

Hi, your patched version does now also find the pppoe interfaces, although I was not yet able to get a picture.

There is a "cleaner" version of the first try at:
<https://github.com/ViToni/igmpproxy/tree/getifaddrs>

In my tests I couldn't get the multicast part to work yet but could detect all the IFs.
Right now I don't know if it's a IGMV2 vs. IMGpV3 or a general issue.
(maybe as in <https://github.com/pali/igmpproxy/issues/12>)

The patch from <https://github.com/ironbits/pfsense-tools/blob/master/pfPorts/igmpproxy/files/patch-fbsd>
contains also some IGMP enhancements which are not included yet.

@pfSense:

Could you provide the source code (or location of it) used to build the igmpproxy for 2.2.6?
It seems there are some changes compared to the sourceforge / github/pali version...

#43 - 05/31/2016 03:45 AM - Stefan Heck

@Victor Toni

I did some Debugging today and finally got your Code working

in mroute-api.c function addMRoute the value of

```
Dp->InVif
```

seems to be wrong. Most likely due to the changed enumeration of interfaces you have coded
I have now assigned a constant value of 1 and the Stream starts to play - at least on my system

```
struct mfctl CtlReq;
int rc;

CtlReq.mfcc_origin    = Dp->OriginAdr;
CtlReq.mfcc_mcastgrp = Dp->McAdr;
CtlReq.mfcc_parent    = 1; /*Dp->InVif*/
```

#44 - 05/31/2016 06:35 AM - Victor Toni

Good catch!

While your are on it could you try to replace line <https://github.com/ViToni/igmpproxy/blob/getifaddrs/src/rtable.c#L270> from

```
newroute->upstrVif = -1;
```

to

```
newroute->upstrVif = ifx;
```

and keep the

```
CtlReq.mfcc_parent = Dp->InVif;
```

#45 - 05/31/2016 07:21 AM - Stefan Heck

Victor,

the Change you have suggested does not work for me.

I also have the Problem that the streams stop responding after a while (< 3 minutes) as soon I'm running different streams on the 4 Entertain Receiver.

I'm investigating this, but already found the changing the value (from 125) in igmpproxy.h

```
#define INTERVAL_QUERY 80
```

dramaticly reduces the Problem

if you like, you can send me a personal message, so we can start communicating more privatly and probably in german ;) It ist definitely my Intention to fix this issues ;)

#46 - 06/01/2016 03:25 AM - Philipp Resch

Stefan Heck wrote:

@Victor Toni

I did some Debugging today and finally got your Code working

in mroute-api.c function addMRoute the value of

[...]

seams to be wrong. Most likly due to the changed enumeration of interfaces you have coded

I have now assigned a constant value of 1 and the Stream starts to play - at least on my system

[...]

Works too for me, hardcoding the value to 1.

Please keep me in the loop, I'm more than willing to test any fixes with the new Telekom BNG infrastructe (no VLAN 8)

EDIT: After having a program running for longer than 12 seconds, the image stops, log file shows this (even with your proposed fix from post [#44](#)):

```
Adding MFC: 87.141.215.251 -> 232.0.10.35, InpVlf: 1
RECV V3 member report from 192.168.1.222 to 224.0.0.22
Updated route entry for 232.0.10.35 on VIF #0
Adding MFC: 87.141.215.251 -> 232.0.10.35, InpVlf: 1
Removing MFC: 87.141.215.251 -> 232.0.10.35, InpVlf: 30
leaveMcGroup: 232.0.10.35 on pppoe1
Removing MFC: 87.141.215.251 -> 232.0.20.35, InpVlf: 30
leaveMcGroup: 232.0.20.35 on pppoe1
```

Inserted route table entry for 232.0.20.35 on VIF #-1
Inserted route table entry for 232.0.10.35 on VIF #-1
The IGMP message was local multicast. Ignoring.
Removing MFC: 87.141.215.251 -> 232.0.10.35, InpVif: 30
MRT_DEL_MFC; Errno(49): Can't assign requested address
Removing MFC: 87.141.215.251 -> 232.0.20.35, InpVif: 30
MRT_DEL_MFC; Errno(49): Can't assign requested address

It seems as if the "-1" interface ID is still being set somewhere.

#47 - 06/01/2016 03:45 AM - Victor Toni

Updated the code at <https://github.com/ViToni/igmpproxy/tree/getifaddrs>
to get the correct value into Dp->InVif and avoid hardcoded values

The change is at <https://github.com/ViToni/igmpproxy/commit/148e7bd96b88a933d6de0857ca6a55db2c24edd4>
from

```
activateRoute(dst, src, upStreamVif[i]-1);
```

to

```
int vifindex = checkVIF->vifindex;  
...  
activateRoute(dst, src, vifindex);
```

Please test.

#48 - 06/01/2016 03:50 AM - Victor Toni

Philipp Resch wrote:

This might be of interest, if not, please remove.

I am already on the "new" German Telekom platform, where VLAN8 no longer exists.
Therefore I need to create an upstream interface on my pppoe interface.

igmpproxy.conf:

quickleave

phyint pppoe1 upstream ratelimit 0 threshold 1
altnet 193.158.0.0/15
altnet 224.0.0.0/4

```
phyint em1 downstream ratelimit 0 threshold 1
altnet 192.168.1.0/24
```

```
phyint pppoe0 disabled
phyint em0 disabled
phyint gif0 disabled
phyint em2 disabled
```

Does this config work with pfSense 2.2.6 and BNG?

#49 - 06/01/2016 03:57 AM - Victor Toni

Philipp Resch wrote:

```
Inserted route table entry for 232.0.20.35 on VIF #-1
Inserted route table entry for 232.0.10.35 on VIF #-1
...
MRT_DEL_MFC; Errno(49): Can't assign requested address
...
MRT_DEL_MFC; Errno(49): Can't assign requested address

It seems as if the "-1" interface ID is still being set somewhere.
```

The "-1" might be the reason for the Errno(49).

#50 - 06/01/2016 06:17 AM - Philipp Resch

The issue still persists, I now get a picture right away with the new version, but it stops after a few seconds.

Here's the log

```
./igmpproxy -d -vvv /tmp/igmpproxy.conf
Adding VIF, Ix 0 FI 0x0 IP 0x0101a8c0 em1, Threshold: 1, Ratelimit: 0
Adding VIF, Ix 1 FI 0x0 IP 0xe18ee95d pppoe1, Threshold: 1, Ratelimit: 0
joinMcGroup: 224.0.0.2 on em1
joinMcGroup: 224.0.0.22 on em1
RECV V3 member report from 192.168.1.1 to 224.0.0.22
The IGMP message was from myself. Ignoring.
The IGMP message was from myself. Ignoring.
RECV V3 member report from 192.168.1.1 to 224.0.0.22
The IGMP message was from myself. Ignoring.
The IGMP message was from myself. Ignoring.
The IGMP message was local multicast. Ignoring.
```

```
RECV V3 member report from 192.168.1.222 to 224.0.0.22
Inserted route table entry for 239.35.10.4 on VIF #0
joinMcGroup: 239.35.10.4 on pppoe1
Adding MFC: 193.158.35.251 -> 239.35.10.4, InpVIf: 1
RECV V3 member report from 192.168.1.222 to 224.0.0.22
Updated route entry for 239.35.10.4 on VIF #0
Adding MFC: 193.158.35.251 -> 239.35.10.4, InpVIf: 1
RECV V3 member report from 192.168.1.222 to 224.0.0.22
Updated route entry for 239.35.10.4 on VIF #0
Adding MFC: 193.158.35.251 -> 239.35.10.4, InpVIf: 1
Removing MFC: 193.158.35.251 -> 239.35.10.4, InpVIf: 1
leaveMcGroup: 239.35.10.4 on pppoe1
Inserted route table entry for 239.35.10.4 on VIF #-1
RECV V3 member report from 192.168.1.222 to 224.0.0.22
RECV V3 member report from 192.168.1.222 to 224.0.0.22
RECV V3 member report from 192.168.1.222 to 224.0.0.22
^Cselect() failure; Errno(4): Interrupted system call
Got a interrupt signal. Exiting.
Removing MFC: 193.158.35.251 -> 239.35.10.4, InpVIf: 1
MRT_DEL_MFC; Errno(49): Can't assign requested address
leaveMcGroup: 239.35.10.4 on pppoe1
MRT_DROP_MEMBERSHIP failed; Errno(49): Can't assign requested address
All routes removed. Routing table is empty.
Shutdown complete....
```

It was working with 2.2.6, I was on BNG right from the start.

#51 - 06/01/2016 12:24 PM - Stefan Heck

Phillip
You may create a Account on
<https://forum.pfsense.org/index.php?board=6.0>
and send me your Email Adresse als PN to user Parsec

I send you an experimental Version of igmpproxy if you like

#52 - 06/02/2016 08:26 AM - Victor Toni

Philipp Resch wrote:

The issue still persists, I now get a picture right away with the new version, but it stops after a few seconds.

Here's the log

```
./igmpproxy -d -vvv /tmp/igmpproxy.conf
```

Could you please try it again? Please use the most recent version from <https://github.com/ViToni/igmpproxy/tree/getifaddrs>
The debug log should have more information (earlier version had an issue with command line parsing of "-vvv"...)

#53 - 06/05/2016 02:52 PM - Andre Vink

It looks like it has something to do with route aging.

The IGMP join is processed neatly and the route is added to the routing table.

however, the route is removed after the route->ageValue has decreased to 0 the route is removed.

The session stays longer if the robustness is increased.

It seems it's not recognized that the route is still an active route.

#54 - 06/06/2016 02:33 AM - Chris Coleman

Andre Vink wrote:

It looks like it has something to do with route aging.

Having looked at the code and at the logs being produced by my igmpproxy in pfSense 2.3 then it looks like the above is what's happening. Connecting to a channel the stream plays flawlessly for just over 4 minutes. With the default robustness set to 2 then my stream dies after approx 4min 15secs, which is very close to the $2 * 125$ that the `DEFAULT_ROBUSTNESS * INTERVAL_QUERY` would give. As this happens I see a leave request in the logs. A quick channel change and I can reconnect instantly and repeat the process.

I'll gladly test any experimental versions :)

EDIT: Testing the version in the post below didn't help. From the log it appears I'm not suffering with the excessive aging bug. I continue to receive membership queries even after the stream has stopped.

#55 - 06/06/2016 02:49 AM - Stefan Heck

You mix up different issues in this thread.

The robustness of igmpproxy has nothing to do with "not recognize upstream interface".

The version from ViToni (<https://github.com/ViToni/igmpproxy/tree/getifaddrs>)

contains some filters to eliminate local multicast traffic.

There is definitely some sort of race condition in the code of updating the routing table

This version runs since a week without any problems and seems to fix the interface detection bug and increases the stability of the service.

It does not fix the bug with the BNG connection for Entertain (Telekom/Germany)

If somebody likes to try this Version you will find it as zip file on <https://www.skeba.de/Dokumente/igmpproxy.zip>

#56 - 06/06/2016 02:53 AM - Andre Vink

Agree, it's what I see as well.
Looking at the logs you'll see the aging counter decrementing and the route removed when it hits zero.
Flip the channel and the process starts again on this channel.
Just for a check I increased the DEFAULT_ROBUSTNESS to 20, recompiled and checked again. Now the stream plays longer and you'll see the aging counter decreasing from 10 to zero.

#57 - 06/06/2016 03:08 AM - Andre Vink

In my opinion the mixup is a result of IGMPproxy not recognizing vlan and PPPoE interfaces.
The version from ViToni recognizes the vlan interfaces perfectly and hence anyone (including me) having problems with the vlan interfaces check this version.
Now we hit this routing issue and that what's reported here.

#58 - 06/06/2016 03:42 AM - Stefan Heck

@Andre Vink
Did you try the Version I have mentioned above? It is the same as the current Version from ViToni plus eliminating some local IGMP Traffic wich may mix up the Routing Table.

#59 - 06/06/2016 03:51 AM - Andre Vink

I'll check it today.

#60 - 06/06/2016 05:05 AM - Andre Vink

@Stefan Heck
The version you provided shows the same route expiration problem. See the log below.

```
11:56:33,6: Current routing table (Age active routes):
11:56:33,6: -----
11:56:33,6: No routes in table...
11:56:33,6: -----
11:56:36,591: RECV V2 member report from 192.168.1.35 to 224.0.252.152
11:56:36,591: Should insert group 224.0.252.152 (from: 192.168.1.35) to route table. Vif Ix : 1
11:56:36,591: No existing route for 224.0.252.152. Create new.
11:56:36,591: No routes in table. Insert at beginning.
11:56:36,591: Inserted route table entry for 224.0.252.152 on VIF #1
11:56:36,591: Joining group 224.0.252.152 upstream on IF address 10.58.240.81
11:56:36,591: joinMcGroup: 224.0.252.152 on em0_vlan4
11:56:36,592: MRT_ADD_MEMBERSHIP on Vif 0 from 10.58.240.81
11:56:36,592:
11:56:36,592: Current routing table (Insert Route):
11:56:36,592: -----
11:56:36,592: #0: Dst: 224.0.252.152, Age:2, St: I, OutVifs: 0x00000002
11:56:36,592: -----
11:56:36,702: Route activate request from 213.75.167.10 to 224.0.252.152 on VIF[0]
11:56:36,702: Vif bits : 0x00000002
11:56:36,702: Setting TTL for Vif 1 to 1
11:56:36,702: Adding MFC: 213.75.167.10 -> 224.0.252.152, InpVIF: 0
11:56:36,702:
11:56:36,702: Current routing table (Activate Route):
11:56:36,702: -----
11:56:36,702: #0: Src0: 213.75.167.10, Dst: 224.0.252.152, Age:2, St: A, OutVifs: 0x00000002
11:56:36,702: -----
11:56:41,285: About to call timeout 296 (#0)
11:56:41,285: Aging routes in table.
11:56:41,285:
11:56:41,285: Current routing table (Age active routes):
```

```

11:56:41,285: -----
11:56:41,285: #0: Src0: 213.75.167.10, Dst: 224.0.252.152, Age:1, St: A, OutVifs: 0x00000002
11:56:41,285: -----
11:56:41,285: About to call timeout 298 (#1)
11:56:47,4: About to call timeout 297 (#0)
11:56:47,4: SENT Membership query from 192.168.1.1 to 224.0.0.1
11:56:47,4: Sent membership query from 192.168.1.1 to 224.0.0.1. Delay: 10
11:56:47,5: Created timeout 299 (#0) - delay 10 secs
11:56:47,5: (Id:299, Time:10)
11:56:47,5: Created timeout 300 (#1) - delay 115 secs
11:56:47,5: (Id:299, Time:10)
11:56:47,5: (Id:300, Time:115)
11:56:50,5: About to call timeout 299 (#0)
11:56:50,5: Aging routes in table.
11:56:50,5: Removing group 224.0.252.152. Died of old age.
11:56:50,5: Removed route entry for 224.0.252.152 from table.
11:56:50,5: Vif bits : 0x00000002
11:56:50,5: Setting TTL for Vif 1 to 1
11:56:50,6: Removing MFC: 213.75.167.10 -> 224.0.252.152, InpVif: 0
11:56:50,6: Leaving group 224.0.252.152 upstream on IF address 10.58.240.81
11:56:50,6: leaveMcGroup: 224.0.252.152 on em0_vlan4
11:56:50,6: MRT_DROP_MEMBERSHIP on Vif 0 from 10.58.240.81
11:56:50,6:
11:56:50,6: Current routing table (Remove route):
11:56:50,6: -----
11:56:50,6: No routes in table...
11:56:50,6: -----
11:56:50,6:
11:56:50,6: Current routing table (Age active routes):
11:56:50,6: -----
11:56:50,6: No routes in table...
11:56:50,6: -----
11:56:50,7: Route activate request from 213.75.167.10 to 224.0.252.152 on VIF[0]
11:56:50,7: No table entry for 224.0.252.152 [From: 213.75.167.10]. Inserting route.
11:56:50,7: No existing route for 224.0.252.152. Create new.
11:56:50,7: No routes in table. Insert at beginning.
11:56:50,7: Inserted route table entry for 224.0.252.152 on VIF #-1
11:56:50,7: No downstream listeners for group 224.0.252.152. No join sent.
11:56:50,7:
11:56:50,7: Current routing table (Insert Route):
11:56:50,7: -----
11:56:50,7: #0: Dst: 224.0.252.152, Age:2, St: I, OutVifs: 0x00000000
11:56:50,7: -----
11:56:50,7:
11:56:50,7: Current routing table (Activate Route):
11:56:50,7: -----
11:56:50,7: #0: Src0: 213.75.167.10, Dst: 224.0.252.152, Age:2, St: A, OutVifs: 0x00000000
11:56:50,7: -----
11:56:54,902: About to call timeout 300 (#0)
11:56:54,902: SENT Membership query from 192.168.1.1 to 224.0.0.1
11:56:54,902: Sent membership query from 192.168.1.1 to 224.0.0.1. Delay: 10
11:56:54,903: Created timeout 301 (#0) - delay 10 secs
11:56:54,903: (Id:301, Time:10)
11:56:54,903: Created timeout 302 (#1) - delay 115 secs
11:56:54,903: (Id:301, Time:10)
11:56:54,903: (Id:302, Time:115)
11:56:57,903: About to call timeout 301 (#0)
11:56:57,903: Aging routes in table.
11:56:57,903:
11:56:57,903: Current routing table (Age active routes):
11:56:57,903: -----
11:56:57,903: #0: Src0: 213.75.167.10, Dst: 224.0.252.152, Age:1, St: A, OutVifs: 0x00000000
11:56:57,903: -----
11:56:59,831: RECV Membership query from 10.60.137.196 to 224.0.0.1
11:57:02,832: About to call timeout 302 (#0)
11:57:02,832: SENT Membership query from 192.168.1.1 to 224.0.0.1
11:57:02,832: Sent membership query from 192.168.1.1 to 224.0.0.1. Delay: 10
11:57:02,833: Created timeout 303 (#0) - delay 10 secs
11:57:02,833: (Id:303, Time:10)
11:57:02,833: Created timeout 304 (#1) - delay 115 secs
11:57:02,833: (Id:303, Time:10)
11:57:02,833: (Id:304, Time:115)
11:57:07,3: About to call timeout 303 (#0)
11:57:07,3: Aging routes in table.
11:57:07,3: Removing group 224.0.252.152. Died of old age.

```

```
11:57:07,3: Removed route entry for 224.0.252.152 from table.
11:57:07,3: Vif bits : 0x00000000
11:57:07,3: Removing MFC: 213.75.167.10 -> 224.0.252.152, InpVif: 0
11:57:07,3: MRT_DEL_MFC; Errno(49): Can't assign requested address
11:57:07,3:
11:57:07,3: Current routing table (Remove route):
11:57:07,3: -----
11:57:07,3: No routes in table...
11:57:07,3: -----
11:57:07,3:
11:57:07,3: Current routing table (Age active routes):
11:57:07,3: -----
11:57:07,3: No routes in table...
11:57:07,3: -----
```

#61 - 06/06/2016 05:52 AM - Andrew -

It looks to me that, for whatever reason, the replies to the membership query aren't getting back to igmpproxy. Igmpproxy sends out periodic membership queries (which you can see from the log), and any client on the network that wants to continue to receive the multicast traffic (i.e. the set top boxes) should reply to confirm their continuing membership of the multicast group. If igmpproxy doesn't receive that confirmation, it will age-out the membership of the group ... which seems to be what's happening here. If the replies aren't getting through, increasing the timeout only delays the inevitable. The question appears to be why it's not receiving the replies. That could conceivably be related to the VLAN issue which was the original bug identified.

#62 - 06/06/2016 06:34 AM - Stefan Heck

@Andrew
yes indeed

While the IP 192.168.1.35 sends a Membership request to 224.0.252.152 11:56:36,591: Should insert group 224.0.252.152 (from: 192.168.1.35) to route table. Vif Ix : 1
it does not reply to the Query Message at
11:56:47,4: SENT Membership query from 192.168.1.1 to 224.0.0.1

Unfortunately the Messages older than 11:56:33,6 are missing in the log.

When you restart igmpproxy without changing the channel on TV, does the stream restarts automatically or do you need to change channel?

have you allowed "IP Options" passing from your VLAN to the Interface igmpproxy running on?

#63 - 06/06/2016 07:37 AM - Andre Vink

To check the IGMP Membership I made two traces on the firewall.

The first is made with the IGMP daemon I used on version 2.2.6. It is not the default PFSense 2.2.6 daemon. Unfortunately I don't know from where I have this one.

But it shows the membership reports. Check 192.167.1.35. It's sending reports for group 224.0.251.109.

TV is working with this daemon. not perfect but it's ok

```
[2.3.1-RELEASE][admin@firewall.familievink.nl]/root: tcpdump -n -i em1_vlan10 igmp
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on em1_vlan10, link-type EN10MB (Ethernet), capture size 65535 bytes
14:12:37.424331 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:12:37.424484 IP 192.168.1.1 > 224.0.0.1: igmp query v3
14:12:38.030587 IP 192.168.1.5 > 224.0.0.251: igmp v2 report 224.0.0.251
14:12:38.420283 IP 192.168.1.34 > 224.3.2.6: igmp v2 report 224.3.2.6
14:12:38.479067 IP 192.168.1.3 > 224.0.1.60: igmp v2 report 224.0.1.60
14:12:40.446538 IP 192.168.1.5 > 239.255.255.250: igmp v2 report 239.255.255.250
14:12:41.003785 IP 192.168.1.43 > 224.0.0.251: igmp v2 report 224.0.0.251
14:12:41.054336 IP 192.168.1.34 > 239.255.255.250: igmp v2 report 239.255.255.250
14:12:41.578596 IP 192.168.1.3 > 224.0.0.251: igmp v2 report 224.0.0.251
14:12:44.638442 IP 192.168.1.10 > 224.168.168.168: igmp v2 report 224.168.168.168
14:12:44.653209 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:12:44.842129 IP 192.168.1.35 > 224.0.251.109: igmp v2 report 224.0.251.109
14:12:44.963245 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:13:08.853297 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:13:08.853456 IP 192.168.1.1 > 224.0.0.1: igmp query v3
14:13:09.017882 IP 192.168.1.42 > 224.0.0.251: igmp v2 report 224.0.0.251
14:13:09.133074 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:13:10.175606 IP 192.168.1.3 > 224.0.0.251: igmp v2 report 224.0.0.251
14:13:12.062601 IP 192.168.1.5 > 239.255.255.250: igmp v2 report 239.255.255.250
14:13:13.020223 IP 192.168.1.43 > 224.0.0.251: igmp v2 report 224.0.0.251
14:13:13.232730 IP 192.168.1.34 > 239.255.255.250: igmp v2 report 239.255.255.250
14:13:13.992845 IP 192.168.1.10 > 224.168.168.168: igmp v2 report 224.168.168.168
14:13:14.175101 IP 192.168.1.3 > 224.0.1.60: igmp v2 report 224.0.1.60
14:13:14.470640 IP 192.168.1.5 > 224.0.0.251: igmp v2 report 224.0.0.251
14:13:15.363111 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:13:15.546013 IP 192.168.1.35 > 224.0.251.109: igmp v2 report 224.0.251.109
14:13:17.014275 IP 192.168.1.34 > 224.3.2.6: igmp v2 report 224.3.2.6
14:15:13.242502 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:15:13.242750 IP 192.168.1.1 > 224.0.0.1: igmp query v3
14:17:18.433397 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:17:18.433559 IP 192.168.1.1 > 224.0.0.1: igmp query v3
14:17:18.550621 IP 192.168.1.5 > 224.0.0.251: igmp v2 report 224.0.0.251
14:17:19.506246 IP 192.168.1.34 > 239.255.255.250: igmp v2 report 239.255.255.250
14:17:20.081934 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:17:20.447772 IP 192.168.1.3 > 224.0.1.60: igmp v2 report 224.0.1.60
14:17:21.398750 IP 192.168.1.8 > 224.0.0.251: igmp v2 report 224.0.0.251
14:17:22.622234 IP 192.168.1.10 > 224.168.168.168: igmp v2 report 224.168.168.168
14:17:23.082113 IP 192.168.1.35 > 224.0.251.109: igmp v2 report 224.0.251.109
14:17:23.176493 IP 192.168.1.35 > 224.3.2.6: igmp v2 report 224.3.2.6
14:17:24.158248 IP 192.168.1.42 > 224.0.0.251: igmp v2 report 224.0.0.251
14:17:26.521917 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:17:28.161608 IP 192.168.1.43 > 224.0.0.251: igmp v2 report 224.0.0.251
14:17:28.342739 IP 192.168.1.5 > 239.255.255.250: igmp v2 report 239.255.255.250
14:19:23.434067 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:19:23.434350 IP 192.168.1.1 > 224.0.0.1: igmp query v3
14:19:23.962175 IP 192.168.1.10 > 224.168.168.168: igmp v2 report 224.168.168.168
14:19:24.203627 IP 192.168.1.42 > 224.0.0.251: igmp v2 report 224.0.0.251
14:19:24.388156 IP 192.168.1.35 > 224.0.251.109: igmp v2 report 224.0.251.109
14:19:24.534032 IP 192.168.1.3 > 224.0.0.251: igmp v2 report 224.0.0.251
14:19:25.811485 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:19:26.622219 IP 192.168.1.34 > 239.255.255.250: igmp v2 report 239.255.255.250
14:19:27.914149 IP 192.168.1.35 > 224.3.2.6: igmp v2 report 224.3.2.6
14:19:27.946165 IP 192.168.1.35 > 239.255.255.250: igmp v2 report 239.255.255.250
14:19:28.333586 IP 192.168.1.3 > 224.0.1.60: igmp v2 report 224.0.1.60
14:19:30.246567 IP 192.168.1.5 > 239.255.255.250: igmp v2 report 239.255.255.250
14:19:30.758627 IP 192.168.1.5 > 224.0.0.251: igmp v2 report 224.0.0.251
14:19:31.961397 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:19:33.094225 IP 192.168.1.34 > 224.3.2.6: igmp v2 report 224.3.2.6
```

The below trace is made with the daemon provided by Stefan.

It looks quite different.

- no IGMPv3 queries
- IGMPv2 queries also from the secondary IP address that is used for pfBlockerNG/DNSBL
- no membership reports. Actually, no membership reports at all for the 224.0.0.0/8 net. only for 239.x multicast addresses.

```
[2.3.1-RELEASE][admin@firewall.familievink.nl]/root: tcpdump -n -i em1_vlan10 igmp
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on em1_vlan10, link-type EN10MB (Ethernet), capture size 65535 bytes
14:04:12.818343 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:04:12.818493 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:04:13.905618 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:04:21.003460 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:04:21.003659 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:04:21.625624 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:04:26.415594 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:04:31.019570 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:04:31.019789 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:04:34.025512 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:04:35.355554 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:04:37.023473 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:04:37.023738 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:04:40.547609 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:04:43.475433 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:04:47.002410 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:04:47.002583 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:04:52.685457 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:04:53.007317 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:04:53.007407 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:04:54.125462 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:04:57.995361 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:05:01.003359 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:05:01.003482 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:05:03.182293 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:05:08.125391 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:05:10.003321 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:05:10.003460 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:05:11.745243 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:05:16.021375 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:05:16.021508 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:05:18.555431 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:05:20.475409 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:05:27.008411 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:05:27.008720 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:05:29.905299 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:05:31.775373 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:05:33.027453 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:05:33.027631 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:05:35.435275 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:05:39.305258 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:05:41.002460 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:05:41.002719 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:05:46.565200 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:05:48.885101 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:05:50.002459 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:05:50.002713 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:05:53.710848 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:05:54.025135 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:05:56.009994 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:05:56.010118 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:05:56.585135 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
14:06:02.175052 IP 192.168.1.6 > 239.255.255.253: igmp v2 report 239.255.255.253
14:06:07.003389 IP 192.168.1.1 > 224.0.0.1: igmp query v2
14:06:07.003573 IP 10.10.10.1 > 224.0.0.1: igmp query v2
14:06:08.195021 IP 192.168.1.178 > 239.255.255.253: igmp v2 report 239.255.255.253
```

#64 - 06/06/2016 08:33 AM - Andrew -

Just a hunch, but I suspect the second query coming from 10.10.10.1 is part of the problem. It's querying the same multicast address as the first query. I wonder whether the network client (i.e. the set top box) is responding to the 10.10.10.1 query (not the 192.168.1.1 query), meaning the reply gets swallowed up by the virtual IP.

I noticed the following config line in an earlier post:

```
altnet 10.0.0.0/8
```

A couple of things you might try:

(1) If multicast streams are originating from your ISP on that 10.0.0.0/8 address range, you might wish to change the virtual IP used for pfBlockerNG into a different range to avoid conflicts.

(2) Alternatively, if igmpproxy is detecting the virtual IP used for pfBlockerNG when it shouldn't be, try adding a line to your config to disable any interfaces that igmpproxy shouldn't be working on (i.e. so you don't get duplicate membership queries from different addresses). Given the new version of igmpproxy picks up VLANs, there may be additional interfaces you need explicitly to disable, which don't cause a problem for the earlier version where they're not detected at all.

#65 - 06/06/2016 08:35 AM - Stefan Heck

Just to compare with the same Version on my System:

MY router (192.168.1.254) periodically sends a Membership query (224.0.0.1) to the TV (192.168.1.213) and this responds correctly to the Multicastaddress it is requesting (239.35.100.3)
The reply message of 192.168.1.213 keeps the Age of the route at 2

```
14:46:34,963: Current routing table (Insert Route):
14:46:34,963: -----
14:46:34,963: #0: Dst: 224.0.0.2, Age:2, St: I, OutVifs: 0x00000001
14:46:34,963: #1: Dst: 239.35.100.3, Age:2, St: I, OutVifs: 0x00000001
14:46:34,963: -----
14:46:35,102: About to call timeout 1 (#0)
14:46:35,102: Aging routes in table.
14:46:35,102:
14:46:35,102: Current routing table (Age active routes):
14:46:35,102: -----
14:46:35,102: #0: Dst: 224.0.0.2, Age:1, St: I, OutVifs: 0x00000001
14:46:35,102: #1: Dst: 239.35.100.3, Age:2, St: I, OutVifs: 0x00000001
14:46:35,102: -----
14:46:39,481: About to call timeout 2 (#0)
14:46:39,481: SENT Membership query from 192.168.1.254 to 224.0.0.1
14:46:39,481: Sent membership query from 192.168.1.254 to 224.0.0.1. Delay: 10
14:46:39,481: Created timeout 3 (#0) - delay 10 secs
14:46:39,481: (Id:3, Time:10)
14:46:39,481: Created timeout 4 (#1) - delay 21 secs
14:46:39,481: (Id:3, Time:10)
14:46:39,481: (Id:4, Time:21)
14:46:39,481: RECV Membership query from 192.168.1.254 to 224.0.0.1
14:46:42,435: RECV V2 member report from 192.168.1.213 to 239.35.100.3
14:46:42,435: Should insert group 239.35.100.3 (from: 192.168.1.213) to route table. Vif Ix : 0
14:46:42,435: Updated route entry for 239.35.100.3 on VIF #0
14:46:42,435:
14:46:42,435: Current routing table (Insert Route):
14:46:42,435: -----
14:46:42,435: #0: Dst: 224.0.0.2, Age:1, St: I, OutVifs: 0x00000001
14:46:42,435: #1: Dst: 239.35.100.3, Age:2, St: I, OutVifs: 0x00000001
14:46:42,435: -----
```

14:46:42,485: RECV V2 member report from 192.168.1.209 to 239.35.100.3
14:46:42,485: Should insert group 239.35.100.3 (from: 192.168.1.209) t

#66 - 06/06/2016 09:23 AM - Andre Vink

To be sure the 10.x address is not messing everything up I removed the address from the interface.

Sadly it didn't help very much

```
16:09:30,4: Current routing table (Age active routes):
16:09:30,4: -----
16:09:30,4: No routes in table...
16:09:30,4: -----
16:09:30,4: Route activate request from 213.75.167.10 to 224.0.251.109 on VIF[0]
16:09:30,4: No table entry for 224.0.251.109 [From: 213.75.167.10]. Inserting route.
16:09:30,4: No existing route for 224.0.251.109. Create new.
16:09:30,4: No routes in table. Insert at beginning.
16:09:30,4: Inserted route table entry for 224.0.251.109 on VIF #-1
16:09:30,4: No downstream listeners for group 224.0.251.109. No join sent.
16:09:30,4:
16:09:30,4: Current routing table (Insert Route):
16:09:30,4: -----
16:09:30,5: #0: Dst: 224.0.251.109, Age:2, St: I, OutVifs: 0x00000000
16:09:30,5: -----
16:09:30,5:
16:09:30,5: Current routing table (Activate Route):
16:09:30,5: -----
16:09:30,5: #0: Src0: 213.75.167.10, Dst: 224.0.251.109, Age:2, St: A, OutVifs: 0x00000000
16:09:30,5: -----
16:09:33,435: About to call timeout 29 (#0)
16:09:33,436: SENT Membership query from 192.168.1.1 to 224.0.0.1
16:09:33,436: Sent membership query from 192.168.1.1 to 224.0.0.1. Delay: 10
16:09:33,436: Created timeout 30 (#0) - delay 10 secs
16:09:33,436: (Id:30, Time:10)
16:09:33,436: Created timeout 31 (#1) - delay 115 secs
16:09:33,436: (Id:30, Time:10)
16:09:33,436: (Id:31, Time:115)
16:09:36,438: About to call timeout 30 (#0)
16:09:36,438: Aging routes in table.
16:09:36,438:
16:09:36,438: Current routing table (Age active routes):
16:09:36,438: -----
16:09:36,438: #0: Src0: 213.75.167.10, Dst: 224.0.251.109, Age:1, St: A, OutVifs: 0x00000000
16:09:36,438: -----
16:09:41,2: About to call timeout 31 (#0)
16:09:41,2: SENT Membership query from 192.168.1.1 to 224.0.0.1
16:09:41,2: Sent membership query from 192.168.1.1 to 224.0.0.1. Delay: 10
16:09:41,2: Created timeout 32 (#0) - delay 10 secs
16:09:41,2: (Id:32, Time:10)
16:09:41,3: Created timeout 33 (#1) - delay 115 secs
16:09:41,3: (Id:32, Time:10)
16:09:41,3: (Id:33, Time:115)
16:09:47,3: About to call timeout 32 (#0)
16:09:47,3: Aging routes in table.
16:09:47,3: Removing group 224.0.251.109. Died of old age.
16:09:47,3: Removed route entry for 224.0.251.109 from table.
16:09:47,3: Vif bits : 0x00000000
16:09:47,3: Removing MFC: 213.75.167.10 -> 224.0.251.109, InpVif: 0
16:09:47,3: MRT_DEL_MFC; Errno(49): Can't assign requested address
16:09:47,3:
16:09:47,3: Current routing table (Remove route):
16:09:47,3: -----
16:09:47,3: No routes in table...
16:09:47,3: -----
16:09:47,3:
16:09:47,3: Current routing table (Age active routes):
16:09:47,3: -----
16:09:47,3: No routes in table...
16:09:47,3: -----
16:09:50,9: About to call timeout 33 (#0)
16:09:50,9: SENT Membership query from 192.168.1.1 to 224.0.0.1
16:09:50,9: Sent membership query from 192.168.1.1 to 224.0.0.1. Delay: 10
16:09:50,9: Created timeout 34 (#0) - delay 10 secs
16:09:50,9: (Id:34, Time:10)
```

```
16:09:50,9: Created timeout 35 (#1) - delay 115 secs
16:09:50,9: (Id:34, Time:10)
16:09:50,9: (Id:35, Time:115)
16:09:53,11: About to call timeout 34 (#0)
16:09:53,11: Aging routes in table.
```

#67 - 06/13/2016 10:05 AM - J.B. BERLIN

Hello,

I tested the Version from ViToni that Stefan Heck posted in the ZIP.
And for me with old Telekom infrastructure Vlan7(Internet) and Vlan8(IPTV) all is now perfect.

I have the put IPTV Downstream in a seperate Vlan44 /28 Network to the TV's.

Here is the output of working igmpproxy in 2.3.1u1

```
./igmpproxy -d -vvv /tmp/igmpproxy.conf
17:01:43,172: Searching for config file at '/tmp/igmpproxy.conf'
17:01:43,173: Config: Quick leave mode enabled.
17:01:43,173: Config: Got a phyint token.
17:01:43,173: Config: IF: Config for interface igb2_vlan8.
17:01:43,173: Config: IF: Got upstream token.
17:01:43,173: Config: IF: Got ratelimit token '0'.
17:01:43,173: Config: IF: Got threshold token '1'.
17:01:43,173: Config: IF: Got altnet token 239.0.0.0/8.
17:01:43,173: Config: IF: Altnet: Parsed altnet to 239/8.
17:01:43,173: Config: IF: Got altnet token 217.0.119.194/16.
17:01:43,173: Config: IF: Altnet: Parsed altnet to 217.0/16.
17:01:43,173: Config: IF: Got altnet token 193.158.35.0/24.
17:01:43,173: Config: IF: Altnet: Parsed altnet to 193.158.35/24.
17:01:43,173: Config: IF name      : igb2_vlan8
17:01:43,173: Config:   Next ptr      : 0
17:01:43,173: Config:   Ratelimit      : 0
17:01:43,173: Config:   Threshold      : 1
17:01:43,173: Config:   State          : 1
17:01:43,173: Config:   Allowednet ptr : 1017050
17:01:43,173: Config: Got a phyint token.
17:01:43,173: Config: IF: Config for interface igb0_vlan44.
17:01:43,173: Config: IF: Got downstream token.
17:01:43,173: Config: IF: Got ratelimit token '0'.
17:01:43,173: Config: IF: Got threshold token '1'.
17:01:43,173: Config: IF: Got altnet token 192.168.44.0/28.
17:01:43,173: Config: IF: Altnet: Parsed altnet to 192.168.44.0/28.
17:01:43,173: Config: IF name      : igb0_vlan44
17:01:43,173: Config:   Next ptr      : 0
17:01:43,173: Config:   Ratelimit      : 0
17:01:43,173: Config:   Threshold      : 1
17:01:43,173: Config:   State          : 2
17:01:43,173: Config:   Allowednet ptr : 1017090
17:01:43,173: Config: Got a phyint token.
17:01:43,173: Config: IF: Config for interface pppoel.
17:01:43,173: Config: IF: Got disabled token.
17:01:43,173: Config: IF name      : pppoel
17:01:43,173: Config:   Next ptr      : 0
17:01:43,173: Config:   Ratelimit      : 0
17:01:43,173: Config:   Threshold      : 1
17:01:43,173: Config:   State          : 0
17:01:43,173: Config:   Allowednet ptr : 0
17:01:43,173: Config: Got a phyint token.
17:01:43,173: Config: IF: Config for interface igb0.
17:01:43,173: Config: IF: Got disabled token.
17:01:43,173: Config: IF name      : igb0
17:01:43,173: Config:   Next ptr      : 0
17:01:43,173: Config:   Ratelimit      : 0
17:01:43,173: Config:   Threshold      : 1
17:01:43,173: Config:   State          : 0
17:01:43,173: Config:   Allowednet ptr : 0
```

17:01:43,173: Config: Got a phyint token.
17:01:43,173: Config: IF: Config for interface pppoe0.
17:01:43,173: Config: IF: Got disabled token.
17:01:43,174: Config: IF name : pppoe0
17:01:43,174: Config: Next ptr : 0
17:01:43,174: Config: Ratelimit : 0
17:01:43,174: Config: Threshold : 1
17:01:43,174: Config: State : 0
17:01:43,174: Config: Allowednet ptr : 0
17:01:43,174: Config: Got a phyint token.
17:01:43,174: Config: IF: Config for interface igb3.
17:01:43,174: Config: IF: Got disabled token.
17:01:43,174: Config: IF name : igb3
17:01:43,174: Config: Next ptr : 0
17:01:43,174: Config: Ratelimit : 0
17:01:43,174: Config: Threshold : 1
17:01:43,174: Config: State : 0
17:01:43,174: Config: Allowednet ptr : 0
17:01:43,174: Config: Got a phyint token.
17:01:43,174: Config: IF: Config for interface igb0_vlan21.
17:01:43,174: Config: IF: Got disabled token.
17:01:43,174: Config: IF name : igb0_vlan21
17:01:43,174: Config: Next ptr : 0
17:01:43,174: Config: Ratelimit : 0
17:01:43,174: Config: Threshold : 1
17:01:43,174: Config: State : 0
17:01:43,174: Config: Allowednet ptr : 0
17:01:43,174: Config: Got a phyint token.
17:01:43,174: Config: IF: Config for interface igb0_vlan22.
17:01:43,174: Config: IF: Got disabled token.
17:01:43,174: Config: IF name : igb0_vlan22
17:01:43,174: Config: Next ptr : 0
17:01:43,174: Config: Ratelimit : 0
17:01:43,174: Config: Threshold : 1
17:01:43,174: Config: State : 0
17:01:43,174: Config: Allowednet ptr : 0
17:01:43,174: Config: Got a phyint token.
17:01:43,174: Config: IF: Config for interface igb1.
17:01:43,174: Config: IF: Got disabled token.
17:01:43,174: Config: IF name : igb1
17:01:43,174: Config: Next ptr : 0
17:01:43,174: Config: Ratelimit : 0
17:01:43,174: Config: Threshold : 1
17:01:43,174: Config: State : 0
17:01:43,174: Config: Allowednet ptr : 0
17:01:43,174: buildIfVc: Starting...
17:01:43,174: buildIfVc: Interface is non-IP: igb0, sa_family: AF_UNKNOWN (18)
17:01:43,174: buildIfVc: Interface is non-IP: igb0, sa_family: AF_INET6 (28)
17:01:43,175: buildIfVc: Interface igb0 Addr: 192.168.115.1, Flags: 0xffff8843, Network: 192.168.115/24
17:01:43,175: buildIfVc: Interface is non-IP: igb1, sa_family: AF_UNKNOWN (18)
17:01:43,175: buildIfVc: Interface is non-IP: igb1, sa_family: AF_INET6 (28)
17:01:43,175: buildIfVc: Interface igb1 Addr: 192.168.1.20, Flags: 0xffff8843, Network: 192.168.1/24
17:01:43,175: buildIfVc: Interface is non-IP: igb2, sa_family: AF_UNKNOWN (18)
17:01:43,175: buildIfVc: Interface is non-IP: igb2, sa_family: AF_INET6 (28)
17:01:43,175: buildIfVc: Interface is non-IP: igb3, sa_family: AF_UNKNOWN (18)
17:01:43,175: buildIfVc: Interface is non-IP: igb3, sa_family: AF_INET6 (28)
17:01:43,175: buildIfVc: Interface igb3 Addr: 192.168.15.1, Flags: 0xffff8843, Network: 192.168.15/24
17:01:43,175: buildIfVc: Interface is non-IP: pflog0, sa_family: AF_UNKNOWN (18)
17:01:43,175: buildIfVc: Interface is non-IP: pfsync0, sa_family: AF_UNKNOWN (18)
17:01:43,175: buildIfVc: Interface is non-IP: enc0, sa_family: AF_UNKNOWN (18)
17:01:43,175: buildIfVc: Interface is non-IP: lo0, sa_family: AF_UNKNOWN (18)
17:01:43,175: buildIfVc: Interface lo0 Addr: 127.0.0.1, Flags: 0xffff8049, Network: 127/8
17:01:43,175: buildIfVc: Interface is non-IP: lo0, sa_family: AF_INET6 (28)
17:01:43,175: buildIfVc: Interface is non-IP: lo0, sa_family: AF_INET6 (28)
17:01:43,175: buildIfVc: Interface is non-IP: igb2_vlan7, sa_family: AF_UNKNOWN (18)
17:01:43,175: buildIfVc: Interface is non-IP: igb2_vlan7, sa_family: AF_INET6 (28)
17:01:43,175: buildIfVc: Interface is non-IP: igb2_vlan8, sa_family: AF_UNKNOWN (18)
17:01:43,175: buildIfVc: Interface is non-IP: igb2_vlan8, sa_family: AF_INET6 (28)
17:01:43,175: buildIfVc: Interface igb2_vlan8 Addr: 10.216.205.108, Flags: 0xffff8843, Network: 10.216.192/18
17:01:43,175: buildIfVc: Interface is non-IP: igb1_vlan7, sa_family: AF_UNKNOWN (18)
17:01:43,175: buildIfVc: Interface is non-IP: igb1_vlan7, sa_family: AF_INET6 (28)
17:01:43,175: buildIfVc: Interface is non-IP: igb0_vlan44, sa_family: AF_UNKNOWN (18)
17:01:43,175: buildIfVc: Interface is non-IP: igb0_vlan44, sa_family: AF_INET6 (28)
17:01:43,175: buildIfVc: Interface igb0_vlan44 Addr: 192.168.44.1, Flags: 0xffff8843, Network: 192.168.44.0/28
17:01:43,175: buildIfVc: Interface is non-IP: igb0_vlan21, sa_family: AF_UNKNOWN (18)

17:01:43,175: buildIfVc: Interface is non-IP: igb0_vlan21, sa_family: AF_INET6 (28)
17:01:43,175: buildIfVc: Interface igb0_vlan21 Addr: 192.168.21.1, Flags: 0xffff8843, Network: 192.168.21/24
17:01:43,175: buildIfVc: Interface is non-IP: igb0_vlan22, sa_family: AF_UNKNOWN (18)
17:01:43,175: buildIfVc: Interface is non-IP: igb0_vlan22, sa_family: AF_INET6 (28)
17:01:43,175: buildIfVc: Interface igb0_vlan22 Addr: xxx.xxx.xxx.xxx, Flags: 0xffff8843, Network: xxx.xxx.xxx
.xxx/29
17:01:43,175: buildIfVc: Interface is non-IP: pppoe1, sa_family: AF_UNKNOWN (18)
17:01:43,175: buildIfVc: Interface pppoe1 Addr: xxx.xxx.xxx.xxx, Flags: 0xffff88d1, Network: xxx.xxx.xxx.xxx/3
2
17:01:43,175: buildIfVc: Interface is non-IP: pppoe1, sa_family: AF_INET6 (28)
17:01:43,175: buildIfVc: Interface is non-IP: pppoe0, sa_family: AF_UNKNOWN (18)
17:01:43,175: buildIfVc: Interface pppoe0 Addr: xxx.xxx.xxx.xxx, Flags: 0xffff88d1, Network: xxx.xxx.xxx.xxx/3
2
17:01:43,175: buildIfVc: Interface is non-IP: pppoe0, sa_family: AF_INET6 (28)
17:01:43,175: buildIfVc: Too many interfaces, skipping ovpnsl
17:01:43,175: Found config for igb0
17:01:43,175: Found config for igb1
17:01:43,175: Found config for igb3
17:01:43,175: Found config for igb2_vlan8
17:01:43,175: Found config for igb0_vlan44
17:01:43,176: Found config for igb0_vlan21
17:01:43,176: Found config for igb0_vlan22
17:01:43,176: Found config for pppoe1
17:01:43,176: Found config for pppoe0
17:01:43,176: Found upstream IF #0, will assign as upstream Vif 22
17:01:43,176: Adding VIF, Ix 0 Fl 0x0 IP 0x6ccdd80a igb2_vlan8, Threshold: 1, Ratelimit: 0
17:01:43,176: Network for [igb2_vlan8] : 10.216.192/18
17:01:43,176: Network for [igb2_vlan8] : 239/8
17:01:43,176: Network for [igb2_vlan8] : 217.0/16
17:01:43,176: Network for [igb2_vlan8] : 193.158.35/24
17:01:43,176: Adding VIF, Ix 1 Fl 0x0 IP 0x012ca8c0 igb0_vlan44, Threshold: 1, Ratelimit: 0
17:01:43,176: Network for [igb0_vlan44] : 192.168.44.0/28
17:01:43,176: Network for [igb0_vlan44] : 192.168.44.0/28
17:01:43,176: Got 262144 byte buffer size in 0 iterations
17:01:43,176: Joining all-routers group 224.0.0.2 on vif 192.168.44.1
17:01:43,176: joinMcGroup: 224.0.0.2 on igb0_vlan44
17:01:43,176: MRT_ADD_MEMBERSHIP on Vif 1 from 192.168.44.1
17:01:43,176: Joining all igmpv3 multicast routers group 224.0.0.22 on vif 192.168.44.1
17:01:43,176: joinMcGroup: 224.0.0.22 on igb0_vlan44
17:01:43,176: MRT_ADD_MEMBERSHIP on Vif 1 from 192.168.44.1
17:01:43,176: SENT Membership query from 192.168.44.1 to 224.0.0.1
17:01:43,176: Sent membership query from 192.168.44.1 to 224.0.0.1. Delay: 10
17:01:43,176: Created timeout 1 (#0) - delay 10 secs
17:01:43,177: (Id:1, Time:10)
17:01:43,177: Created timeout 2 (#1) - delay 21 secs
17:01:43,177: (Id:1, Time:10)
17:01:43,177: (Id:2, Time:21)
17:01:43,177: RECV V2 member report from 192.168.44.1 to 224.0.0.2
17:01:43,177: The IGMP message was from myself. Ignoring.
17:01:43,177: RECV Membership query from 192.168.44.1 to 224.0.0.1
17:01:44,923: RECV V2 member report from 192.168.44.1 to 224.0.0.2
17:01:44,923: The IGMP message was from myself. Ignoring.
17:01:45,91: RECV Membership query from 10.216.255.254 to 224.0.0.1
17:01:45,734: RECV V2 member report from 192.168.44.2 to 239.35.100.9
17:01:45,734: Should insert group 239.35.100.9 (from: 192.168.44.2) to route table. Vif Ix : 1

```

17:01:45,734: No existing route for 239.35.100.9. Create new.
17:01:45,734: No routes in table. Insert at beginning.
17:01:45,734: Inserted route table entry for 239.35.100.9 on VIF #1
17:01:45,734: Joining group 239.35.100.9 upstream on IF address 10.216.205.108
17:01:45,734: joinMcGroup: 239.35.100.9 on igb2_vlan8
17:01:45,734: MRT_ADD_MEMBERSHIP on Vif 0 from 10.216.205.108
17:01:45,734:
17:01:45,734: Current routing table (Insert Route):
17:01:45,734: -----
17:01:45,734: #0: Dst: 239.35.100.9, Age:2, St: I, OutVifs: 0x00000002
17:01:45,734: -----
17:01:46,166: Route activate request from 193.158.35.177 to 239.35.100.9 on VIF[0]
17:01:46,166: Vif bits : 0x00000002
17:01:46,166: Setting TTL for Vif 1 to 1
17:01:46,166: Adding MFC: 193.158.35.177 -> 239.35.100.9, InpVIF: 0
17:01:46,166:
17:01:46,166: Current routing table (Activate Route):
17:01:46,166: -----
17:01:46,166: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:01:46,167: -----
17:01:46,831: RECV V2 member report from 192.168.44.10 to 239.35.100.9
17:01:46,831: Should insert group 239.35.100.9 (from: 192.168.44.10) to route table. Vif Ix : 1
17:01:46,831: Updated route entry for 239.35.100.9 on VIF #1
17:01:46,831: Vif bits : 0x00000002
17:01:46,831: Setting TTL for Vif 1 to 1
17:01:46,831: Adding MFC: 193.158.35.177 -> 239.35.100.9, InpVIF: 0
17:01:46,831:
17:01:46,831: Current routing table (Insert Route):
17:01:46,831: -----
17:01:46,831: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:01:46,831: -----
17:01:47,11: RECV V2 member report from 192.168.44.11 to 239.35.100.9
17:01:47,11: Should insert group 239.35.100.9 (from: 192.168.44.11) to route table. Vif Ix : 1
17:01:47,11: Updated route entry for 239.35.100.9 on VIF #1
17:01:47,11: Vif bits : 0x00000002
17:01:47,11: Setting TTL for Vif 1 to 1
17:01:47,11: Adding MFC: 193.158.35.177 -> 239.35.100.9, InpVIF: 0
17:01:47,11:
17:01:47,11: Current routing table (Insert Route):
17:01:47,11: -----
17:01:47,11: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:01:47,11: -----
17:01:47,871: RECV V2 member report from 192.168.44.7 to 239.35.100.9
17:01:47,871: Should insert group 239.35.100.9 (from: 192.168.44.7) to route table. Vif Ix : 1
17:01:47,871: Updated route entry for 239.35.100.9 on VIF #1
17:01:47,871: Vif bits : 0x00000002
17:01:47,871: Setting TTL for Vif 1 to 1
17:01:47,871: Adding MFC: 193.158.35.177 -> 239.35.100.9, InpVIF: 0
17:01:47,871:
17:01:47,871: Current routing table (Insert Route):
17:01:47,871: -----
17:01:47,871: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:01:47,871: -----
17:01:48,723: RECV V2 member report from 192.168.44.5 to 239.35.100.9
17:01:48,723: Should insert group 239.35.100.9 (from: 192.168.44.5) to route table. Vif Ix : 1
17:01:48,723: Updated route entry for 239.35.100.9 on VIF #1
17:01:48,723: Vif bits : 0x00000002
17:01:48,724: Setting TTL for Vif 1 to 1
17:01:48,724: Adding MFC: 193.158.35.177 -> 239.35.100.9, InpVIF: 0
17:01:48,724:
17:01:48,724: Current routing table (Insert Route):
17:01:48,724: -----
17:01:48,724: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:01:48,724: -----
17:01:49,872: RECV V2 member report from 192.168.44.7 to 239.35.20.10
17:01:49,872: Should insert group 239.35.20.10 (from: 192.168.44.7) to route table. Vif Ix : 1
17:01:49,872: No existing route for 239.35.20.10. Create new.
17:01:49,873: Found existing routes. Find insert location.
17:01:49,873: Inserting after route 239.35.100.9
17:01:49,873: Inserted route table entry for 239.35.20.10 on VIF #1
17:01:49,873: Joining group 239.35.20.10 upstream on IF address 10.216.205.108
17:01:49,873: joinMcGroup: 239.35.20.10 on igb2_vlan8
17:01:49,873: MRT_ADD_MEMBERSHIP on Vif 0 from 10.216.205.108
17:01:49,873:
17:01:49,873: Current routing table (Insert Route):

```

```

17:01:49,873: -----
17:01:49,873: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:01:49,873: #1: Dst: 239.35.20.10, Age:2, St: I, OutVifs: 0x00000002
17:01:49,873: -----
17:01:49,952: Route activate request from 193.158.35.251 to 239.35.20.10 on VIF[0]
17:01:49,952: Vif bits : 0x00000002
17:01:49,952: Setting TTL for Vif 1 to 1
17:01:49,952: Adding MFC: 193.158.35.251 -> 239.35.20.10, InpVif: 0
17:01:49,952:
17:01:49,952: Current routing table (Activate Route):
17:01:49,952: -----
17:01:49,952: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:01:49,952: #1: Src0: 193.158.35.251, Dst: 239.35.20.10, Age:2, St: A, OutVifs: 0x00000002
17:01:49,952: -----
17:01:50,798: RECV V2 member report from 192.168.44.6 to 239.35.100.9
17:01:50,798: Should insert group 239.35.100.9 (from: 192.168.44.6) to route table. Vif Ix : 1
17:01:50,798: Updated route entry for 239.35.100.9 on VIF #1
17:01:50,798: Vif bits : 0x00000002
17:01:50,798: Setting TTL for Vif 1 to 1
17:01:50,798: Adding MFC: 193.158.35.177 -> 239.35.100.9, InpVif: 0
17:01:50,798:
17:01:50,798: Current routing table (Insert Route):
17:01:50,798: -----
17:01:50,798: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:01:50,798: #1: Src0: 193.158.35.251, Dst: 239.35.20.10, Age:2, St: A, OutVifs: 0x00000002
17:01:50,798: -----
17:01:51,815: RECV V2 member report from 192.168.44.4 to 239.35.100.9
17:01:51,815: Should insert group 239.35.100.9 (from: 192.168.44.4) to route table. Vif Ix : 1
17:01:51,815: Updated route entry for 239.35.100.9 on VIF #1
17:01:51,815: Vif bits : 0x00000002
17:01:51,815: Setting TTL for Vif 1 to 1
17:01:51,815: Adding MFC: 193.158.35.177 -> 239.35.100.9, InpVif: 0
17:01:51,815:
17:01:51,815: Current routing table (Insert Route):
17:01:51,815: -----
17:01:51,815: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:01:51,815: #1: Src0: 193.158.35.251, Dst: 239.35.20.10, Age:2, St: A, OutVifs: 0x00000002
17:01:51,816: -----
17:01:53,812: About to call timeout 1 (#0)
17:01:53,812: Aging routes in table.
17:01:53,812:
17:01:53,812: Current routing table (Age active routes):
17:01:53,812: -----
17:01:53,812: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:01:53,812: #1: Src0: 193.158.35.251, Dst: 239.35.20.10, Age:1, St: A, OutVifs: 0x00000002
17:01:53,812: -----
17:01:56,814: About to call timeout 2 (#0)
17:01:56,814: SENT Membership query from 192.168.44.1 to 224.0.0.1
17:01:56,814: Sent membership query from 192.168.44.1 to 224.0.0.1. Delay: 10
17:01:56,814: Created timeout 3 (#0) - delay 10 secs
17:01:56,814: (Id:3, Time:10)
17:01:56,814: Created timeout 4 (#1) - delay 21 secs
17:01:56,814: (Id:3, Time:10)
17:01:56,814: (Id:4, Time:21)
17:01:56,814: RECV Membership query from 192.168.44.1 to 224.0.0.1
17:01:58,739: RECV V2 member report from 192.168.44.5 to 239.35.100.9
17:01:58,739: Should insert group 239.35.100.9 (from: 192.168.44.5) to route table. Vif Ix : 1
17:01:58,739: Updated route entry for 239.35.100.9 on VIF #1
17:01:58,739: Vif bits : 0x00000002
17:01:58,739: Setting TTL for Vif 1 to 1
17:01:58,739: Adding MFC: 193.158.35.177 -> 239.35.100.9, InpVif: 0
17:01:58,739:
17:01:58,739: Current routing table (Insert Route):
17:01:58,739: -----
17:01:58,739: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:01:58,739: #1: Src0: 193.158.35.251, Dst: 239.35.20.10, Age:1, St: A, OutVifs: 0x00000002
17:01:58,739: -----
17:01:59,378: RECV V2 member report from 192.168.44.7 to 239.35.20.10
17:01:59,378: Should insert group 239.35.20.10 (from: 192.168.44.7) to route table. Vif Ix : 1
17:01:59,379: Updated route entry for 239.35.20.10 on VIF #1
17:01:59,379: Vif bits : 0x00000002
17:01:59,379: Setting TTL for Vif 1 to 1
17:01:59,379: Adding MFC: 193.158.35.251 -> 239.35.20.10, InpVif: 0
17:01:59,379:
17:01:59,379: Current routing table (Insert Route):

```

```

17:01:59,379: -----
17:01:59,379: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:01:59,379: #1: Src0: 193.158.35.251, Dst: 239.35.20.10, Age:1, St: A, OutVifs: 0x00000002
17:01:59,379: -----
17:01:59,742: RECV V2 member report from 192.168.44.2 to 239.35.100.9
17:01:59,742: Should insert group 239.35.100.9 (from: 192.168.44.2) to route table. Vif Ix : 1
17:01:59,742: Updated route entry for 239.35.100.9 on VIF #1
17:01:59,742: Vif bits : 0x00000002
17:01:59,742: Setting TTL for Vif 1 to 1
17:01:59,742: Adding MFC: 193.158.35.177 -> 239.35.100.9, InpVif: 0
17:01:59,742:
17:01:59,742: Current routing table (Insert Route):
17:01:59,742: -----
17:01:59,742: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:01:59,742: #1: Src0: 193.158.35.251, Dst: 239.35.20.10, Age:1, St: A, OutVifs: 0x00000002
17:01:59,742: -----
17:02:01,330: RECV V2 member report from 192.168.44.4 to 239.35.100.9
17:02:01,330: Should insert group 239.35.100.9 (from: 192.168.44.4) to route table. Vif Ix : 1
17:02:01,330: Updated route entry for 239.35.100.9 on VIF #1
17:02:01,330: Vif bits : 0x00000002
17:02:01,330: Setting TTL for Vif 1 to 1
17:02:01,330: Adding MFC: 193.158.35.177 -> 239.35.100.9, InpVif: 0
17:02:01,330:
17:02:01,330: Current routing table (Insert Route):
17:02:01,330: -----
17:02:01,330: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:02:01,330: #1: Src0: 193.158.35.251, Dst: 239.35.20.10, Age:1, St: A, OutVifs: 0x00000002
17:02:01,330: -----
17:02:01,381: RECV V2 member report from 192.168.44.7 to 239.35.100.9
17:02:01,381: Should insert group 239.35.100.9 (from: 192.168.44.7) to route table. Vif Ix : 1
17:02:01,381: Updated route entry for 239.35.100.9 on VIF #1
17:02:01,381: Vif bits : 0x00000002
17:02:01,381: Setting TTL for Vif 1 to 1
17:02:01,381: Adding MFC: 193.158.35.177 -> 239.35.100.9, InpVif: 0
17:02:01,381:
17:02:01,381: Current routing table (Insert Route):
17:02:01,381: -----
17:02:01,381: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:02:01,381: #1: Src0: 193.158.35.251, Dst: 239.35.20.10, Age:1, St: A, OutVifs: 0x00000002
17:02:01,381: -----
17:02:02,342: RECV V2 member report from 192.168.44.10 to 239.35.100.9
17:02:02,342: Should insert group 239.35.100.9 (from: 192.168.44.10) to route table. Vif Ix : 1
17:02:02,342: Updated route entry for 239.35.100.9 on VIF #1
17:02:02,342: Vif bits : 0x00000002
17:02:02,342: Setting TTL for Vif 1 to 1
17:02:02,342: Adding MFC: 193.158.35.177 -> 239.35.100.9, InpVif: 0
17:02:02,342:
17:02:02,342: Current routing table (Insert Route):
17:02:02,342: -----
17:02:02,342: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:02:02,342: #1: Src0: 193.158.35.251, Dst: 239.35.20.10, Age:1, St: A, OutVifs: 0x00000002
17:02:02,342: -----
17:02:02,805: RECV V2 member report from 192.168.44.6 to 239.35.100.9
17:02:02,805: Should insert group 239.35.100.9 (from: 192.168.44.6) to route table. Vif Ix : 1
17:02:02,805: Updated route entry for 239.35.100.9 on VIF #1
17:02:02,805: Vif bits : 0x00000002
17:02:02,805: Setting TTL for Vif 1 to 1
17:02:02,805: Adding MFC: 193.158.35.177 -> 239.35.100.9, InpVif: 0
17:02:02,805:
17:02:02,805: Current routing table (Insert Route):
17:02:02,805: -----
17:02:02,805: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:02:02,805: #1: Src0: 193.158.35.251, Dst: 239.35.20.10, Age:1, St: A, OutVifs: 0x00000002
17:02:02,805: -----
17:02:05,528: RECV V2 member report from 192.168.44.1 to 224.0.0.2
17:02:05,528: The IGMP message was from myself. Ignoring.
17:02:06,522: RECV V2 member report from 192.168.44.11 to 239.35.100.9
17:02:06,522: Should insert group 239.35.100.9 (from: 192.168.44.11) to route table. Vif Ix : 1
17:02:06,522: Updated route entry for 239.35.100.9 on VIF #1
17:02:06,522: Vif bits : 0x00000002
17:02:06,522: Setting TTL for Vif 1 to 1
17:02:06,522: Adding MFC: 193.158.35.177 -> 239.35.100.9, InpVif: 0
17:02:06,522:
17:02:06,522: Current routing table (Insert Route):
17:02:06,522: -----

```

```
17:02:06,522: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:02:06,522: #1: Src0: 193.158.35.251, Dst: 239.35.20.10, Age:1, St: A, OutVifs: 0x00000002
17:02:06,522: -----
17:02:06,522: About to call timeout 3 (#0)
17:02:06,522: Aging routes in table.
17:02:06,522:
17:02:06,522: Current routing table (Age active routes):
17:02:06,522: -----
17:02:06,523: #0: Src0: 193.158.35.177, Dst: 239.35.100.9, Age:2, St: A, OutVifs: 0x00000002
17:02:06,523: #1: Src0: 193.158.35.251, Dst: 239.35.20.10, Age:2, St: A, OutVifs: 0x00000002
17:02:06,523: -----
^C17:02:07,948: select() failure; Errno(4): Interrupted system call
17:02:07,948: Got a interupt signal. Exiting.
17:02:07,948: clean handler called
17:02:07,948: Removing route entry for 239.35.100.9
17:02:07,948: Vif bits : 0x00000002
17:02:07,948: Setting TTL for Vif 1 to 1
17:02:07,948: Removing MFC: 193.158.35.177 -> 239.35.100.9, InpVif: 0
17:02:07,948: Leaving group 239.35.100.9 upstream on IF address 10.216.205.108
17:02:07,948: leaveMcGroup: 239.35.100.9 on igb2_vlan8
17:02:07,948: MRT_DROP_MEMBERSHIP on Vif 0 from 10.216.205.108
17:02:07,948: Removing route entry for 239.35.20.10
17:02:07,948: Vif bits : 0x00000002
17:02:07,948: Setting TTL for Vif 1 to 1
17:02:07,948: Removing MFC: 193.158.35.251 -> 239.35.20.10, InpVif: 0
17:02:07,948: Leaving group 239.35.20.10 upstream on IF address 10.216.205.108
17:02:07,948: leaveMcGroup: 239.35.20.10 on igb2_vlan8
17:02:07,949: MRT_DROP_MEMBERSHIP on Vif 0 from 10.216.205.108
17:02:07,949: All routes removed. Routing table is empty.
17:02:07,949: Shutdown complete....
```

If someone like, I can run some Tests.

Greetings

J.B.

#68 - 06/14/2016 01:59 AM - Victor Toni

There were some issue with the logging on startup in some corner cases. Hopefully these are fixed in:
<https://github.com/ViToni/igmpproxy/tree/logging>

This version should work with IGMPv2 (which I cannot test myself because I am on IGMPv3).
IGMPv3 is not expected to work at that time.

Please do some testing.

(Currently DEVEL_LOGGING is active for better debugging. It could be turned off at a later time, but is actually only important when using the "-d" switch.)

#69 - 06/14/2016 10:47 AM - Greg Myran

This version should work with IGMPv2 (which I cannot test myself because I am on IGMPv3).
IGMPv3 is not expected to work at that time.

Please do some testing.

IGMPv2 working for me with this latest version. Thanks for all your time and effort Victor (and others).

#70 - 06/15/2016 05:46 PM - Victor Toni

Greg Myran wrote:

IGMPv2 working for me with this latest version. Thanks for all your time and effort Victor (and others).

Great! You're welcome. Thanks for testing.

There might have been some issues regarding delVif /addVof which should be fixed with the latest version from:
<https://github.com/ViToni/igmpproxy/tree/logging>

Could you please give it one more try to avoid regressions?

#71 - 06/17/2016 05:16 AM - Prince Adam

@Victor Toni

I'm not quite sure if this might be relevant. But it seems like there is a repository at Github attempting to implement IGMPv3:
<https://github.com/diagprov/igmpproxy>

Do you mind having a look? Maybe you could join your efforts then.

#72 - 06/18/2016 10:57 AM - Greg Myran

Victor Toni wrote:

Could you please give it one more try to avoid regressions?

This latest version has been working and stable for 2+ days now.

I will note that I did have problems in one of my 2 setups (not new to this version). If I enable ipv6 on the downstream interface (using 6RD for IPv6 addressing) the multicast stream starts, runs for about 2 seconds then stops.

I bumped the LogLevel to TRACE and here are the logs with and without IPv6 enabled on the downstream interface

• IPv6 enabled : not working

```
[Trace] 10:17:39,849 log_IP():434: Internet Protocol, Src Addr: [redacted], Dst Addr: [redacted]
[Trace] 10:17:39,849 log_IP():435: Version: 4
[Trace] 10:17:39,849 log_IP():436: Header length: 20
[Trace] 10:17:39,849 log_IP():437: Total Length: 19461
[Trace] 10:17:39,849 log_IP():438: Identification: 0x0000 (0)
[Trace] 10:17:39,849 log_IP():439: Fragment offset: 0
[Trace] 10:17:39,849 log_IP():440: Time to live: 1
[Trace] 10:17:39,849 log_IP():441: Protocol: IP6 hop-by-hop options (0x00)
[Trace] 10:17:39,849 log_IP():443: Header Checksum: 0x8701
[Trace] 10:17:39,849 log_received_IGMP():264: IGMP: received packet
[Trace] 10:17:39,849 log_received_IGMP():267: - received packet from [redacted] shorter (20 bytes) than hd
r+data length (20+19441)
[Debug] 10:17:39,849 acceptIcmp():150: Route activate request from [redacted] to [redacted] on VIF[1]
[Debug] 10:17:39,849 logRouteTable():708:
[Debug] 10:17:39,849 logRouteTable():709: Current routing table (Activate Route):
[Debug] 10:17:39,849 logRouteTable():710: -----
[Debug] 10:17:39,849 logRouteTable():730: #0: Src0: [redacted], Dst: [redacted], Age:2, St: A, OutVifs: 0x0000
0000
[Debug] 10:17:39,849 logRouteTable():730: #1: Src0: [redacted], Dst: [redacted], Age:1, St: A, OutVifs: 0x0000
0000
[Debug] 10:17:39,849 logRouteTable():730: #2: Src0: [redacted], Dst: [redacted], Age:2, St: A, OutVifs: 0x0000
0000
[Debug] 10:17:39,849 logRouteTable():738: -----
```

• IPv6 disabled : working

```
[Trace] 10:21:23,275 log_IP():434: Internet Protocol, Src Addr: [redacted], Dst Addr: [redacted]
[Trace] 10:21:23,275 log_IP():435: Version: 4
[Trace] 10:21:23,275 log_IP():436: Header length: 20
[Trace] 10:21:23,275 log_IP():437: Total Length: 36352
[Trace] 10:21:23,275 log_IP():438: Identification: 0x5c39 (23609)
[Trace] 10:21:23,275 log_IP():439: Fragment offset: 0
[Trace] 10:21:23,275 log_IP():440: Time to live: 1
[Trace] 10:21:23,275 log_IP():441: Protocol: IP6 hop-by-hop options (0x00)
[Trace] 10:21:23,275 log_IP():443: Header Checksum: 0x4401
[Trace] 10:21:23,275 log_received_IGMP():264: IGMP: received packet
[Trace] 10:21:23,275 log_received_IGMP():267: - received packet from [redacted] shorter (20 bytes) than hd
r+data length (20+36332)
[Debug] 10:21:23,275 acceptIcmp():150: Route activate request from [redacted] to [redacted] on VIF[1]
[Debug] 10:21:23,275 internUpdateKernelRoute():673: Vif bits : 0x00000001
[Debug] 10:21:23,275 internUpdateKernelRoute():682: Setting TTL for Vif 0 to 1
[Debug] 10:21:23,275 addMRoute():245: Adding MFC (MRT_ADD_MFC): [redacted] -> [redacted], InpVIf: 1
[Debug] 10:21:23,275 logRouteTable():708:
```

```
[Debug] 10:21:23,275 logRouteTable():709: Current routing table (Activate Route):
[Debug] 10:21:23,275 logRouteTable():710: -----
[Debug] 10:21:23,275 logRouteTable():730: #0: Src0: [redacted], Dst: [redacted], Age:2, St: A, OutVifs: 0x0000
0001
[Debug] 10:21:23,275 logRouteTable():730: #1: Src0: [redacted], Dst: [redacted], Age:2, St: A, OutVifs: 0x0000
0001
[Debug] 10:21:23,275 logRouteTable():730: #2: Src0: [redacted], Dst: [redacted], Age:2, St: A, OutVifs: 0x0000
0001
[Debug] 10:21:23,275 logRouteTable():738: -----
```

It looks like in the IPv6 enabled case the igmp packet is being identified as "Version 4" but still sees "IP6 hop-by-hop options" so something is incorrectly parsing that packet and causing it to skip the 'internUpdateKernelRoute' process.

I'm guessing this is a VERY rare use case, but it may be worth noting anyway.

#73 - 06/18/2016 05:22 PM - Victor Toni

It is not supposed to work! igmpproxy supports only IPv4 / IGMP.

Regarding multicast IGMP is for "pure" IPv4, MLD is for IPv6.
Usually there is no mix between.

There are projects like <https://github.com/mcproxy/mcproxy> (currently linux only) which support both but still only on the same protocol version (IPv4 <=> IPv4, IPv6 <=> IPv6).

There is another project (<https://unfix.org/projects/ecmh/>) which does the translation between IPv4/IGMP <=> IPv6/MLD but more as routing so the proxy part is missing as I understand the project.

#74 - 06/18/2016 09:07 PM - Greg Myran

I understand igmpproxy only supports IPv4 / IGMP. In both cases above the STB is IPv4 only and speaks IGMPv2. The difference is the interface on the PFSense box is also IPv4 only in the working case and dual stack IPv4 and IPv6 in the broken case. I am now starting to suspect the new method of finding network interfaces may see the additional address family and marks the interface as "non-IP" and ignores the IPv4 address it had previously found, almost like it should find the IPv4 address and break out of a loop instead of continuing to look for additional addresses??

```
[Debug] 10:10:52,212 buildIfVc():418: buildIfVc: Interface igb3 Addr: 10.0.2.1, Flags: 0xffff8843, Network: 0.0.0.0/32
[Debug] 10:10:52,212 buildIfVc():405: buildIfVc: Interface is non-IP: igb3, sa_family: AF_INET6 (28)
[Debug] 10:10:52,212 buildIfVc():405: buildIfVc: Interface is non-IP: igb3, sa_family: AF_INET6 (28)
```

I will play around with this more tomorrow.

Also, FWIW, the version of igmpproxy that came with PFSense pre 2.3 works perfectly with the dual stack PFSense interface. It wasn't until I installed the latest working builds that I had to disable IPv6 on the PFSense side.

#75 - 06/19/2016 08:41 AM - Greg Myran

I did some more troubleshooting and here are some additional details.

If I start igmpproxy with IPv6 disabled, then enable IPv6, igmpproxy continues to work fine. If I then restart igmpproxy, it no longer works properly. With it running in the not working state I disabled IPv6 on the interface and it continues to not work until I again restart with only IPv4 enabled.

I then started capturing packets on the wire to see what was happening on the LAN segment. What I found is that when igmpproxy starts without IPv6 enabled I see IGMPv2 messages (which is what my multicast feeds require), if I start with IPv6 enabled I see IGMPv3 messages except the occasional IGMPv2 membership query from PFSense itself.

I reverted back to the igmpproxy binary from PFSense 2.2 and did the same tests/captures and in both the IPv4 only and dual stack cases I see IGMPv2 messages (and it works in both cases).

#76 - 06/20/2016 11:53 AM - Victor Toni

Greg Myran wrote:

[...] I am now starting to suspect the new method of finding network interfaces may see the additional address family and marks the interface as "non-IP" and ignores the IPv4 address it had previously found, almost like it should find the IPv4 address and break out of a loop instead of continuing to look for additional addresses??

```
[Debug] 10:10:52,212 buildIfVc():418: buildIfVc: Interface igb3 Addr: 10.0.2.1, Flags: 0xffff8843, Network: 0.0.0.0/32
```

```
[Debug] 10:10:52,212 buildIfVc():405: buildIfVc: Interface is non-IP: igb3, sa_family: AF_INET6 (28)
```

```
[Debug] 10:10:52,212 buildIfVc():405: buildIfVc: Interface is non-IP: igb3, sa_family: AF_INET6 (28)
```

This log snippet says that 3 IFs were found, one IPv4 and two IPv6. IFs without an address and IFs which are not AF_INET (IPv4) are ignored / skipped. After "buildIfVc()" there should be some "addVIF()" entries showing which devices will be used for multicast.

Also, FWIW, the version of igmpproxy that came with PFSense pre 2.3 works perfectly with the dual stack PFSense interface. It wasn't until I installed the latest working builds that I had to disable IPv6 on the PFSense side.

The igmpproxy from pfSense 2.2.6 is a custom one for which I don't have access to the source code so I cannot really compare what's happening. I tried to use the patch from <https://github.com/ironbits/pfsense-tools/blob/master/pfPorts/igmpproxy/files/patch-fbsd> but had only unsatisfactory results.

#77 - 06/20/2016 05:02 PM - Greg Myran

I want to thank you again for your efforts. I've moved my multicast traffic onto it's own physical interface on my PFSense box so I can keep dual stack for LAN (PC/phone/other) clients and have only IPv4 on the STB interface. I'm completely willing to continue to assist if you want to troubleshoot this further but I also understand it's probably a low priority issue and for all I know this issue is unique to my STB's and multicast service.

#78 - 06/22/2016 02:04 PM - Victor Toni

Greg Myran wrote:

I want to thank you again for your efforts.

You're welcome,

I've moved my multicast traffic onto it's own physical interface on my PFSense box so I can keep dual stack for LAN (PC/phone/other) clients and have only IPv4 on the STB interface. I'm completely willing to continue to assist if you want to troubleshoot this further but I also understand it's probably a low priority issue and for all I know this issue is unique to my STB's and multicast service.

I have a similar setup but only for IPv4. Unfortunately my provider is using IGMPv3 which I cannot get to work.

#79 - 06/22/2016 02:14 PM - Victor Toni

As said my provider is using IGMPv3 which I cannot get to work (so I will stay on 2.2.6 for the time being).

I've invested some time to adapt the patch from <https://github.com/ironbits/pfsense-tools/blob/master/pfPorts/igmpproxy/files/patch-fbsd> to the seemingly functional version (for IGMPv2, the logging branch: <https://github.com/ViToni/igmpproxy/tree/logging>). Even after working around the quirks introduced by the multi-upstream / dynamic IF reload patch and with the sending of IGMPv3 enabled there seems to be some magic sauce missing.

As I have to shift priorities until the end of the year there won't be any progress from my side until that time. So feel free to dig in the current state of the code and do some hacking ;-)

<https://github.com/ViToni/igmpproxy/tree/IGMPv3>

Even if I have no time for hacking myself I would be more than happy to accept pull requests if someone happens to find the "magic sauce" :-)

Happy coding!

#80 - 07/10/2016 03:53 AM - Chris Buechler

- Target version changed from 2.3.2 to 2.4.0

#81 - 08/28/2016 02:32 PM - Lars Veldcholte

What is the status on this? I believe I have the same issue (pfSense 2.3).

All my interfaces are VLANs. I just moved my set-top box to a separate VLAN because I suspect problems with IGMP snooping on my switch.

My STB used to be in VLAN 1, now it is in VLAN 5.

My new igmpproxy config is:

```
quickleave
phyint vmx0_vlan4 upstream ratelimit 0 threshold 1
altnet 213.75.167.0/24

phyint vmx0_vlan5 downstream ratelimit 0 threshold 1

phyint vmx0_vlan34 disabled
phyint vmx0_vlan1 disabled
phyint vmx0_vlan2 disabled
phyint gif0 disabled
```

The problem is that the new downstream interface (vmx0_vlan5) is not detected by igmpproxy. If I run igmpproxy verbose, I see:

```
buildIfVc: Interface lo0 Addr: 127.0.0.1, Flags: 0xffff8049, Network: 127/8
buildIfVc: Interface vmx0_vlan1 Addr: 192.168.2.1, Flags: 0xffff8843, Network: 192.168.2/24
buildIfVc: Interface vmx0_vlan2 Addr: 192.168.3.1, Flags: 0xffff8843, Network: 192.168.3/24
buildIfVc: Interface vmx0_vlan4 Addr: 10.10.12.186, Flags: 0xffff8843, Network: 10.10.12/22
buildIfVc: Interface vmx0_vlan34 Addr: 84.245.30.184, Flags: 0xffff8843, Network: 84.245.30/24
```

But vmx0_vlan5 is missing.

I don't understand why it does find the other interfaces, but not vmx0_vlan5. I've checked everything, the configuration is basically identical. Is there anything I might have missed that causes igmpproxy to skip the interface?

@Victor Toni: Would upgrading to your fork of igmpproxy solve things? If so, how do I install it?

I downloaded the source to my pfSense box and installed *autoconf* and *automake*, but the configure script complains that there is no C compiler found. I guess I still have to install the compiler? How do I do that in pfSense?

Also, which branch should I be using? There's quite a lot of them in your repo ;)

EDIT:

I couldn't get a build environment working on my pfSense box, so I compiled it on another FreeBSD machine.

Initially I picked your newest branch (*IGMPv3*). That seemed to work (detected the interface, no errors), but I still had no signal on my TV. So I tried again with the *getifaddrs* branch. This one actually works!

#82 - 08/31/2016 06:22 AM - Victor Toni

Lars Veldcholte wrote:

What is the status on this? I believe I have the same issue (pfSense 2.3).

[...]

I don't understand why it does find the other interfaces, but not vmx0_vlan5. I've checked everything, the configuration is basically identical. Is there anything I might have missed that causes igmpproxy to skip the interface?

It's the way the interfaces are looked up (internally) so its not directly related to your config.

@Victor Toni: Would upgrading to your fork of igmpproxy solve things? If so, how do I install it?

I downloaded the source to my pfSense box and installed *autoconf* and *automake*, but the configure script complains that there is no C compiler found. I guess I still have to install the compiler? How do I do that in pfSense?

Also, which branch should I be using? There's quite a lot of them in your repo ;)

EDIT:

I couldn't get a build environment working on my pfSense box, so I compiled it on another FreeBSD machine.

Initially I picked your newest branch (*IGMPv3*). That seemed to work (detected the interface, no errors), but I still had no signal on my TV. So I tried again with the *getifaddrs* branch. This one actually works!

I'm working with VirtualBox so that I can test multiple OS & versions. (Sometimes there are great differences between needed headers...)

Could you please test the *logging* branch? This one is supposed to work with IGMPv2.

(And if the *getifaddrs* works for you so should *logging* but with enhanced logging and some minor fixes)

<https://github.com/ViToni/igmpproxy/tree/logging>

#83 - 09/02/2016 11:06 AM - Lars Veldcholte

Victor Toni wrote:

It's the way the interfaces are looked up (internally) so its not directly related to your config.

I get that, but there must be something that differentiates the interfaces that are detected from the interfaces that aren't. It can't be pure chance, right?

What I meant is that I couldn't find any property of the interface that wasn't detected that could cause it to be not detected. It was basically identical in all ways I could think of to the interface that are detected.

I'm working with VirtualBox so that I can test multiple OS & versions. (Sometimes there are great differences between needed headers...)

Could you please test the *logging* branch? This one is supposed to work with IGMPv2.

(And if the *getifaddrs* works for you so should *logging* but with enhanced logging and some minor fixes)

<https://github.com/ViToni/igmpproxy/tree/logging>

I just compiled and installed the *logging* branch, this one works as well. Well, just as well.

I still have problems with TV, but that might not be related to igmpproxy at all. Like I said in my first post in this thread, I moved my STB to a separate VLAN because I suspected issues with IGMP snooping on my switch. Well... it turned out that wasn't it, because my problem isn't solved.

This problem is that the IPTV stream cuts out after 10-15 minutes. I have no idea what causes it. I don't see any suspicious messages in the igmpproxy log when it happens. igmpproxy still receives membership reports from the STB, and the IPTV route is still maintained in igmpproxy's routing table.

#84 - 09/03/2016 04:53 PM - Jorge M. Oliveira

(I'm using original version of igmpproxy without any changes)

There is one thing I find very interesting.

On my test configuration, IGMP Proxy recognizes the first 4 configured VLANs and the 5th is not.

If I run ifconfig on console, the first 4 VLANs have:

```
flags=8a43<UP,BROADCAST,RUNNING,ALLMULTI,SIMPLEX,MULTICAST>
```

while the 5th has:

flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST>

Maybe ALLMULTI missing is the reason for this?

#85 - 09/03/2016 06:48 PM - Jorge M. Oliveira

By the way, I've coded a very hackish workaround (for version 2.3.3) that one can execute via Diagnostics > Command Prompt > Execute PHP Commands

```
if (!is_array($config['igmpproxy']['igmpentry']) || count($config['igmpproxy']['igmpentry']) == 0) {
    return;
}
if (!is_array($config['vlans']['vlan']) || count($config['vlans']['vlan']) == 0) {
    return;
}
$igmprealifs = $igmpvlans = $othervlans = $vlans = array();
foreach ($config['igmpproxy']['igmpentry'] as $igmpentry) {
    $igmprealifs[] = get_real_interface($igmpentry['ifname']);
}
foreach ($config['vlans']['vlan'] as $vlan) {
    if (in_array($vlan['vlanif'], $igmprealifs)) {
        $igmpvlans[] = $vlan;
    } else {
        $othervlans[] = $vlan;
    }
}
$vlans = array_merge($igmpvlans, $othervlans);
$config['vlans']['vlan'] = $vlans;
write_config();
```

This will reorder the system VLANs so IGMP ones are available first. System must be rebooted to apply changes. Proceed at your own risk and remember to backup configuration before trying this! I've tested it a little and worked here.

#86 - 09/04/2016 03:03 PM - Jorge M. Oliveira

- File patch-src__os-freebsd.h added

- File patch-src_ifvc.c added

- File patch-src_igmpproxy.c added

Some brand new patches for the version that ships on freebsd-ports:

patch-src_os-freebsd.h is based on <https://github.com/pali/igmpproxy/commit/feaeed153535fa8e6cdb201886b90d1a8eabad67> adding support to freebsd 11 vs the default version on freebsd-ports. (edit: seems the devel does already have this, must have gotten the original from the RELENG_2_3 tree).

patch-src_ifvc.c is based on freebsd-ports version plus custom tweaks to be sure all vifs (up to a reasonable limit, at least the first 15 if lucky enough) are able to be listed.

patch-src_igmpproxy.c is partially based on <https://github.com/pali/igmpproxy/commit/85e240727305b156097ee7aa0f0c4473a136291f> default behaviour is still that an interface is IF_STATE_DOWNSTREAM but if an interface is explicitly disabled, it won't count towards addVIF() neither vifcount.

The source for this problem is that maximum interfaces is 40 but FreeBSD routines return 3 to 4 different descriptors for the same interface(AF_LINK, AF_INET6 (with or without linklocal) and AF_INET) so you'll hit the 40 limit with just 13 to 10 interfaces. FreeBSD also has some "dummy" interfaces for internal tasks, those eat space as well. My patches come as a baid-aid for this issue and I'd appreciate developers taking a look and merging them to the tree (PR being prepared).

Some logs after patch is applied (regarding the interfaces "mess")

```
Interface is not AF_INET, skipping em0 (family 18)
Interface is not AF_INET, skipping em0 (family 28)
buildIfVc: Interface em0 Addr: 192.168.111.158, Flags: 0xffff8843, Network: 192.168.111/24
Interface is not AF_INET, skipping pflog0 (family 18)
Interface is not AF_INET, skipping pfsync0 (family 18)
Interface is not AF_INET, skipping enc0 (family 18)
Interface is not AF_INET, skipping lo0 (family 18)
buildIfVc: Interface lo0 Addr: 127.0.0.1, Flags: 0xffff8049, Network: 127/8
Interface is not AF_INET, skipping lo0 (family 28)
Interface is not AF_INET, skipping lo0 (family 28)
Interface is not AF_INET, skipping em0_vlan1 (family 18)
Interface is not AF_INET, skipping em0_vlan1 (family 28)
buildIfVc: Interface em0_vlan1 Addr: 192.168.0.1, Flags: 0xffff8843, Network: 192.168.0/24
Interface is not AF_INET, skipping em0_vlan1 (family 28)
Interface is not AF_INET, skipping em0_vlan2 (family 18)
Interface is not AF_INET, skipping em0_vlan2 (family 28)
buildIfVc: Interface em0_vlan2 Addr: 192.168.1.1, Flags: 0xffff8843, Network: 192.168.1/24
Interface is not AF_INET, skipping em0_vlan3 (family 18)
Interface is not AF_INET, skipping em0_vlan3 (family 28)
buildIfVc: Interface em0_vlan3 Addr: 192.168.2.1, Flags: 0xffff8843, Network: 192.168.2/24
Interface is not AF_INET, skipping em0_vlan4 (family 18)
Interface is not AF_INET, skipping em0_vlan4 (family 28)
buildIfVc: Interface em0_vlan4 Addr: 192.168.3.1, Flags: 0xffff8843, Network: 192.168.3/24
```

Tested on pfSense 2.3.3 and 2.4.0!

PR: <https://github.com/pfsense/FreeBSD-ports/pull/182>

^ Has an extra commit that has a rewrite to use getifaddrs(), be sure to take a look at it! (The patch-src_ifvc.c attached here is now outdated)

#87 - 09/05/2016 04:49 AM - Jorge M. Oliveira

- File *igmproxy_all.zip* added

#88 - 09/05/2016 04:51 AM - Jorge M. Oliveira

EDIT: *igmproxy_all.zip* is somewhat good, except no IGMPv3 support. The version on pfSense 2.2 supported those packets.

I've added in my post above custom *igmproxy* builds based on FreeBSD-ports PR [#182](#) for pfSense 2.3 (amd64 and i386) and 2.4 (amd64). Have fun and good luck! :)

#89 - 09/05/2016 12:33 PM - Jorge M. Oliveira

- File *igmproxy_20160905_1818.zip* added

EDIT: This version is bugged. Please use the previous. kthxbye.

New version based on the reviewed patch. I believe my work in this area is complete (for now).

I am hoping 2.3.3 will soon include so no more custom patching/replacing system files.

Maybe in the future this has further room to improve and 2.4.0 will see an improved *igmproxy* with much more support and options.

#90 - 09/05/2016 01:06 PM - Victor Toni

Jorge M. Oliveira wrote:

New version based on the reviewed patch. I believe my work in this area is complete (for now).

I am hoping 2.3.3 will soon include so no more custom patching/replacing system files.

Maybe in the future this has further room to improve and 2.4.0 will see an improved *igmproxy* with much more support and options.

Does your version support IGMPv3?

#91 - 09/05/2016 03:24 PM - Jorge M. Oliveira

No, my purpose was only to fix this long standing bug and (attempt) to make it work again as it should.

I leave the implementation details of IGMPv3 to your and others knowledge and we'll see what happens when a new *igmproxy* version is officially released.

#92 - 09/05/2016 03:57 PM - Victor Toni

Jorge M. Oliveira wrote:

No, my purpose was only to fix this long standing bug and (attempt) to make it work again as it should.

I leave the implementation details of IGMPv3 to your and others knowledge and we'll see what happens when a new *igmproxy* version is officially released.

This version does solve the interfaces issue already and some other minor bugs, too:
<https://github.com/ViToni/igmpproxy/tree/logging>
(What it does not solve is the IGMPv3 compatibility issue.)

Is there a special case were this version did not work for you?

#93 - 09/09/2016 01:50 PM - Jorge M. Oliveira

Victor Toni wrote:

This version does solve the interfaces issue already and some other minor bugs, too:
<https://github.com/ViToni/igmpproxy/tree/logging>
(What it does not solve is the IGMPv3 compatibility issue.)

Is there a special case were this version did not work for you?

No, your version appears to work well.

#94 - 09/18/2016 11:24 AM - Dora Paula

Hi,

first of all, thank you very much for your hard work!

Due to missing IPTV (I just ordered it last week) I instead tried to get upnp-multicasts (for dlina) running across two subnets ... without success. Here are the details:

I replaced 2.3.2-igmpproxy-binary (/usr/local/sbin/igmpproxy) with the one offered in the zip file linked here:
<https://redmine.pfsense.org/issues/6099#note-89>

0. Local Network Setup: =====
10.112.100.0/24 (dlina-server: .10)
10.112.5.0/24 (dlina-client: .128)
various other vlans (should not matter here)

1. the configuration file =====

```
[2.3.2-RELEASE][admin@pfsensebox]/root: cat /tmp/igmpproxy.conf
```

```
##-----
```

```
1. Enable Quickleave mode (Sends Leave instantly)
```

```
##-----  
quickleave  
phyint igb0_vlan11 upstream ratelimit 0 threshold 1  
altnet 10.112.100.0/24
```

```
phyint igb1_vlan50 downstream ratelimit 0 threshold 1  
altnet 10.112.5.0/24
```

```
phyint pppoe0 disabled  
phyint igb1_vlan20 disabled  
phyint gif0 disabled  
phyint igb0_vlan10 disabled  
phyint igb1_vlan30 disabled  
phyint igb0_vlan15 disabled
```

phyint igb1_vlan40 disabled

2. start the custom igmpproxy-binary via console with debugging-option:

```
=====
Searching for config file at '/tmp/igmpproxy.conf'
Config: Quick leave mode enabled.
Config: Got a phyint token.
Config: IF: Config for interface igb0_vlan11.
Config: IF: Got upstream token.
Config: IF: Got ratelimit token '0'.
Config: IF: Got threshold token '1'.
Config: IF: Got altnet token 10.112.100.0/24.
Config: IF: Altnet: Parsed altnet to 0.112.100/24.
^^^^^^^^^^strange IP address
IF name : igb0_vlan11
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 1
Allowednet ptr : 1017050
Config: Got a phyint token.
Config: IF: Config for interface igb1_vlan50.
Config: IF: Got downstream token.
Config: IF: Got ratelimit token '0'.
Config: IF: Got threshold token '1'.
Config: IF: Got altnet token 10.112.5.0/24.
Config: IF: Altnet: Parsed altnet to 0.112.5/24.
^^^^^^^^^^strange IP address
IF name : igb1_vlan50
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 2
Allowednet ptr : 1017070
Config: Got a phyint token.
Config: IF: Config for interface pppoe0.
Config: IF: Got disabled token.
IF name : pppoe0
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface igb1_vlan20.
Config: IF: Got disabled token.
IF name : igb1_vlan20
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface gif0.
Config: IF: Got disabled token.
IF name : gif0
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface igb0_vlan10.
Config: IF: Got disabled token.
IF name : igb0_vlan10
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface igb1_vlan30.
Config: IF: Got disabled token.
IF name : igb1_vlan30
```

Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface igb0_vlan15.
Config: IF: Got disabled token.
IF name : igb0_vlan15
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface igb1_vlan40.
Config: IF: Got disabled token.
IF name : igb1_vlan40
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
buildfVc: Starting...
buildfVc: Interface is not AF_INET, skipping igb0 (family 18)
buildfVc: Interface is not AF_INET, skipping igb0 (family 28)
buildfVc: Interface is not AF_INET, skipping igb1 (family 18)
buildfVc: Interface is not AF_INET, skipping igb1 (family 28)
buildfVc: Interface is not AF_INET, skipping igb2 (family 18)
buildfVc: Interface is not AF_INET, skipping pflog0 (family 18)
buildfVc: Interface is not AF_INET, skipping pfsync0 (family 18)
buildfVc: Interface is not AF_INET, skipping enc0 (family 18)
buildfVc: Interface is not AF_INET, skipping lo0 (family 18)
buildfVc: Interface lo0 Addr: 127.0.0.1, Flags: 0xffff8049, Network: 127/8
buildfVc: Interface is not AF_INET, skipping lo0 (family 28)
buildfVc: Interface is not AF_INET, skipping lo0 (family 28)
buildfVc: Interface is not AF_INET, skipping igb0_vlan7 (family 18)
buildfVc: Interface is not AF_INET, skipping igb0_vlan7 (family 28)
buildfVc: Interface is not AF_INET, skipping igb0_vlan10 (family 18)
buildfVc: Interface is not AF_INET, skipping igb0_vlan10 (family 28)
buildfVc: Interface igb0_vlan10 Addr: 10.112.1.254, Flags: 0xffff8843, Network: 10.112.1/24
buildfVc: Interface is not AF_INET, skipping igb0_vlan10 (family 28)
buildfVc: Interface is not AF_INET, skipping igb0_vlan11 (family 18)
buildfVc: Interface is not AF_INET, skipping igb0_vlan11 (family 28)
buildfVc: Interface igb0_vlan11 Addr: 10.112.100.254, Flags: 0xffff8843, Network: 10.112.100/24
buildfVc: Interface is not AF_INET, skipping igb0_vlan11 (family 28)
buildfVc: Interface is not AF_INET, skipping igb0_vlan15 (family 18)
buildfVc: Interface is not AF_INET, skipping igb0_vlan15 (family 28)
buildfVc: Interface igb0_vlan15 Addr: 10.112.15.254, Flags: 0xffff8843, Network: 10.112.15/24
buildfVc: Interface is not AF_INET, skipping igb0_vlan15 (family 28)
buildfVc: Interface is not AF_INET, skipping igb1_vlan20 (family 18)
buildfVc: Interface is not AF_INET, skipping igb1_vlan20 (family 28)
buildfVc: Interface igb1_vlan20 Addr: 10.112.2.254, Flags: 0xffff8843, Network: 10.112.2/24
buildfVc: Interface is not AF_INET, skipping igb1_vlan20 (family 28)
buildfVc: Interface is not AF_INET, skipping igb1_vlan30 (family 18)
buildfVc: Interface is not AF_INET, skipping igb1_vlan30 (family 28)
buildfVc: Interface igb1_vlan30 Addr: 10.112.3.254, Flags: 0xffff8843, Network: 10.112.3/24
buildfVc: Interface is not AF_INET, skipping igb1_vlan30 (family 28)
buildfVc: Interface is not AF_INET, skipping igb1_vlan40 (family 18)
buildfVc: Interface is not AF_INET, skipping igb1_vlan40 (family 28)
buildfVc: Interface igb1_vlan40 Addr: 10.112.4.254, Flags: 0xffff8843, Network: 10.112.4/24
buildfVc: Interface is not AF_INET, skipping igb1_vlan40 (family 28)
buildfVc: Interface is not AF_INET, skipping igb1_vlan50 (family 18)
buildfVc: Interface is not AF_INET, skipping igb1_vlan50 (family 28)
buildfVc: Interface igb1_vlan50 Addr: 10.112.5.254, Flags: 0xffff8943, Network: 10.112.5/24
buildfVc: Interface is not AF_INET, skipping igb1_vlan60 (family 18)
buildfVc: Interface is not AF_INET, skipping igb1_vlan60 (family 28)
buildfVc: Interface is not AF_INET, skipping pppoe0 (family 18)
buildfVc: Interface is not AF_INET, skipping pppoe0 (family 28)
buildfVc: Interface pppoe0 Addr: 12.12.12.12, Flags: 0xffff88d1, Network: 12.12.12.12/32
buildfVc: Interface is not AF_INET, skipping gif0 (family 18)
buildfVc: Interface is not AF_INET, skipping gif0 (family 28)
buildfVc: Interface is not AF_INET, skipping gif0 (family 28)
buildfVc: Interface is not AF_INET, skipping ovpn1 (family 18)
buildfVc: Interface is not AF_INET, skipping ovpn1 (family 28)

```
buildfVc: Interface ovpsn1 Addr: 10.112.128.1, Flags: 0xffff8051, Network: 10.112.128/24
buildfVc: Interface is not AF_INET, skipping ovpsn1 (family 28)
Found config for igb0_vlan10
Found config for igb0_vlan11
Found config for igb0_vlan15
Found config for igb1_vlan20
Found config for igb1_vlan30
Found config for igb1_vlan40
Found config for igb1_vlan50
Found config for pppoe0
adding VIF, Ix 0 FI 0x0 IP 0xfe64700a igb0_vlan11, Threshold: 1, Ratelimit: 0
Network for [igb0_vlan11] : 10.112.100/24
Network for [igb0_vlan11] : 0.112.100/24
^^^^^^^^^^strange IP address
adding VIF, Ix 1 FI 0x0 IP 0xfe05700a igb1_vlan50, Threshold: 1, Ratelimit: 0
Network for [igb1_vlan50] : 10.112.5/24
Network for [igb1_vlan50] : 0.112.5/24
^^^^^^^^^^strange IP address
adding VIF, Ix 2 FI 0x0 IP 0x0180700a ovpsn1, Threshold: 1, Ratelimit: 0
Network for [ovpsn1] : 10.112.128/24
Got 262144 byte buffer size in 0 iterations
Joining all-routers group 224.0.0.2 on vif 10.112.5.254
joinMcGroup: 224.0.0.2 on igb1_vlan50
Joining all-routers group 224.0.0.2 on vif 10.112.128.1
joinMcGroup: 224.0.0.2 on ovpsn1
SENT Membership query from 10.112.5.254 to 224.0.0.1
Sent membership query from 10.112.5.254 to 224.0.0.1. Delay: 10
SENT Membership query from 10.112.128.1 to 224.0.0.1
Sent membership query from 10.112.128.1 to 224.0.0.1. Delay: 10
Created timeout 1 (#0) - delay 10 secs
(Id:1, Time:10)
Created timeout 2 (#1) - delay 21 secs
(Id:1, Time:10)
(Id:2, Time:21)
About to call timeout 1 (#0)
Aging routes in table.
Current routing table (Age active routes):
-----
No routes in table...
-----
```

3. Started the TV (dlna-client multicast) igmpproxy output continues:

=====

```
About to call timeout 2 (#0)
SENT Membership query from 10.112.5.254 to 224.0.0.1
Sent membership query from 10.112.5.254 to 224.0.0.1. Delay: 10
SENT Membership query from 10.112.128.1 to 224.0.0.1
Sent membership query from 10.112.128.1 to 224.0.0.1. Delay: 10
Created timeout 3 (#0) - delay 10 secs
(Id:3, Time:10)
Created timeout 4 (#1) - delay 21 secs
(Id:3, Time:10)
(Id:4, Time:21)
About to call timeout 3 (#0)
Aging routes in table.
Current routing table (Age active routes):
-----
No routes in table...
-----
```

```
The source address 10.112.5.128 for group 239.255.255.250, is not in any valid net for upstream VIF.
The source address 10.112.5.128 for group 239.255.255.250, is not in any valid net for upstream VIF.
The source address 10.112.5.128 for group 239.255.255.250, is not in any valid net for upstream VIF.
```

4. Start dlna-server, igmpproxy-outputs continues: =====

```
Route activate request from 10.112.100.10 to 239.255.255.250
No table entry for 239.255.255.250 [From: 10.112.100.10]. Inserting route.
No existing route for 239.255.255.250. Create new.
No routes in table. Insert at beginning.
Inserted route table entry for 239.255.255.250 on VIF #-1
No downstream listeners for group 239.255.255.250. No join sent.

Current routing table (Insert Route):
-----
```

#0: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000000

Current routing table (Activate Route):

#0: Src: 10.112.100.10, Dst: 239.255.255.250, Age:2, St: A, OutVifs: 0x00000000

About to call timeout 4 (#0)

SENT Membership query from 10.112.5.254 to 224.0.0.1

Sent membership query from 10.112.5.254 to 224.0.0.1. Delay: 10

SENT Membership query from 10.112.128.1 to 224.0.0.1

Sent membership query from 10.112.128.1 to 224.0.0.1. Delay: 10

Created timeout 5 (#0) - delay 10 secs

(Id:5, Time:10)

Created timeout 6 (#1) - delay 115 secs

(Id:5, Time:10)

(Id:6, Time:115)

About to call timeout 5 (#0)

Aging routes in table.

Current routing table (Age active routes):

#0: Src: 10.112.100.10, Dst: 239.255.255.250, Age:1, St: A, OutVifs: 0x00000000

select() failure; Errno(4): Interrupted system call

Got a interrupt signal. Exiting.

clean handler called

Removing route entry for 239.255.255.250

Vif bits : 0x00000000

Removing MFC: 10.112.100.10 -> 239.255.255.250, InpVif: 0

MRT_DEL_MFC; Errno(49): Can't assign requested address

All routes removed. Routing table is empty.

Shutdown complete....

#95 - 10/13/2016 04:36 PM - Philipp Haefelfinger

I have the same issue like Dora Paule with the version: igmpproxy_20160905_1818.zip

There is no such problem with the version posted before: igmpproxy_all.zip

pfsense: 2.3.2_p1

pfsense: 2.3.3.a.20161013.0748

The configuration looks valid to me (Swisscom TV2 in switzerland)

##-----

```
## Enable Quickleave mode (Sends Leave instantly)
##-----
quickleave
phyint igb2_vlan10 upstream ratelimit 0 threshold 1
altnet 224.0.0.0/4
altnet 213.3.72.0/24
altnet 195.186.0.0/16

phyint igb3_vlan1 downstream ratelimit 0 threshold 1
altnet a.b.c.d/xx
```

The output during startup is quite strange. the token is correct but the parsed value is wrong:

Output

```
Allowednet ptr : 1017050
Config: Got a phyint token.
Config: IF: Config for interface igb2_vlan10.
Config: IF: Got upstream token.
Config: IF: Got ratelimit token '0'.
Config: IF: Got threshold token '1'.
Config: IF: Got altnet token 224.0.0.0/4.
Config: IF: Altnet: Parsed altnet to 0/4.
Config: IF: Got altnet token 213.3.72.0/24.
Config: IF: Altnet: Parsed altnet to 0.3.72/24.
Config: IF: Got altnet token 195.186.0.0/16.
Config: IF: Altnet: Parsed altnet to 0.0/16.
```

...

```
adding VIF, Ix 0 F1 0x0 IP 0xf2cb3dbc igb2_vlan10, Threshold: 1, Ratelimit: 0
  Network for [igb2_vlan10] : a.b.c.d/22
  Network for [igb2_vlan10] : 0/4
  Network for [igb2_vlan10] : 0.3.72/24
  Network for [igb2_vlan10] : 0.0/16
```

Looking at this the following error message:

```
The source address a.b.c.d for group e.f.g.h, is not in any valid net for upstream VIF.
```

is right because there is no match.

Somewhere between the "all" version and "20160905_1818" must be a change that causes this bug.

I really like pfsense but this igmp issue drives me crazy. The funny part is it started happening to me after I added two more nics to my pfsense box (total counting 6 NICs now) and had to reassign the interfaces. Before that I had no issue with it. I did not change the name of the vlan nor the assignment names. Even the WAN is on the same physical interface as before but the interfaces have new igb numbers. The two new ports are igb0 and igb1 and all the previous available ports are now incremented by 2. So igb0 became igb2 etc. This seems to be the root cause of the issue.

I hope this helps locating the error and please please fix this thing. As ipTV is getting more popular it would be really nice if this is working properly.

Thanks

Philipp

#96 - 10/17/2016 04:33 AM - Michiel Lowijs

I have the same problem with the 20160905_1818 version.

The _all version works fine on ISP XS4All in The Netherlands with the following config file:

```
quickleave
phyint em0_vlan4 upstream ratelimit 0 threshold 1
altnet 213.75.0.0/16
altnet 10.86.0.0/16
altnet 224.0.0.0/4

phyint em1 downstream ratelimit 0 threshold 1
altnet 192.168.1.0/24

phyint pppoe0 disabled
```

#97 - 11/09/2016 04:41 PM - Chris Becker

Found sth on different site:

[[<https://sourceforge.net/p/igmpproxy/bugs/4/#472a>]]

So for at least with DE-Telekom they must have stopped to supply IPv4 Addresses via DHCP on VLAN8

So the Reason this doesn't work on 2.3.0 is that there is no IPv4 Address on the IPTV Interface:

Added a dump one manual and IGMPProxy starts again ... :

AlgvBh7l.png

Cheers,

Chris

#98 - 11/10/2016 04:04 PM - Philipp Haefelfinger

That's interesting. But unfortunately this is not the case for my system. Swisscom transmits everything on vlan10 and regulates the speed with QoS limits.

I only have IPv4 Addresses because ipv6 currently does not work very well with the swisscom 6rd infrastructure. So this is still an issue.

Cheers,

Philipp

#99 - 12/15/2016 08:20 AM - Harald Gutmann

Dear Maintainers,

@Jorge M. Oliveira, thank you for your work to fix this issues.

I'm using the igmpproxy to control a Sonos System which is located in a different VLAN.
Therefore can confirm that both versions are working nicely with Setup.

igmpproxy_all.zip
igmpproxy_20160905_1818.zip

Kindly included the improved version in the next release of pfSense.

Best regards,
Harald Gutmann

#100 - 12/23/2016 02:41 PM - Luiz Souza

- Status changed from Confirmed to Feedback

- % Done changed from 0 to 100

Fix committed.

Thanks!

#101 - 12/29/2016 03:19 PM - Philipp Haefelfinger

Is this commit in the latest build applied? If yes, there seems to be something buggy with the fix.
I just updated my box to the pfsense 2.4 build 2.4.0.b.20161229.1430.

/var/etc/igmpproxy.conf

```
##-----  
## Enable Quickleave mode (Sends Leave instantly)  
##-----  
quickleave  
phyint igb2_vlan10 upstream ratelimit 0 threshold 1  
altnet 224.0.0.0/4  
altnet 213.3.72.0/24  
altnet 195.186.0.0/16  
  
phyint igb3_vlan1 downstream ratelimit 0 threshold 1  
altnet 192.168.10.0/24  
# some disabled interfaces cut out
```

Using the latest build I get the following type of messages from the proxy:

The source address 192.168.10.XX for group 239.255.255.250, is not in any valid net for upstream VIF.

Starting with "/usr/local/sbin/igmpproxy -d -vv /var/etc/igmpproxy.conf" shows the following registrations:

```
adding VIF, Ix 0 Fl 0x0 IP 0xf2cb3dbc igb2_vlan10, Threshold: 1, Ratelimit: 0
    Network for [igb2_vlan10] : 188.61.200/22
    Network for [igb2_vlan10] : 0/4
    Network for [igb2_vlan10] : 0.3.72/24
    Network for [igb2_vlan10] : 0.0/16
    Network for [igb2_vlan10] : 0.0/16
adding VIF, Ix 1 Fl 0x0 IP 0x010aa8c0 igb3_vlan1, Threshold: 1, Ratelimit: 0
    Network for [igb3_vlan1] : 192.168.10/24
    Network for [igb3_vlan1] : 0.168.10/24
adding VIF, Ix 2 Fl 0x0 IP 0x019da8c0 ovpsn1, Threshold: 1, Ratelimit: 0
    Network for [ovpsn1] : 192.168.157/24
```

This does not seem right. I would expect the following output from igmpproxy:

```
adding VIF, Ix 0 Fl 0x0 IP 0xf2cb3dbc igb2_vlan10, Threshold: 1, Ratelimit: 0
    Network for [igb2_vlan10] : 188.61.200/22
    Network for [igb2_vlan10] : 224/4
    Network for [igb2_vlan10] : 213.3.72/24
    Network for [igb2_vlan10] : 195.186/16
    Network for [igb2_vlan10] : 188.61/16
adding VIF, Ix 1 Fl 0x0 IP 0x010aa8c0 igb3_vlan1, Threshold: 1, Ratelimit: 0
    Network for [igb3_vlan1] : 192.168.10/24
adding VIF, Ix 2 Fl 0x0 IP 0x019da8c0 ovpsn1, Threshold: 1, Ratelimit: 0
    Network for [ovpsn1] : 192.168.157/24
```

Hope this helps to find the cause of the invalid parsing.

Greets

Philipp

#102 - 12/31/2016 05:04 AM - Vincent Gijzen

Philipp Haefelfinger wrote:

Is this commit in the latest build applied? If yes, there seems to be something buggy with the fix.
I just updated my box to the pfsense 2.4 build 2.4.0.b.20161229.1430.

```
/var/etc/igmpproxy.conf  
[...]
```

Using the latest build I get the following type of messages from the proxy:

```
[...]
```

Starting with "/usr/local/sbin/igmpproxy -d -vv /var/etc/igmpproxy.conf" shows the following registrations:
[...]

This does not seem right. I would expect the following output from igmpproxy:
[...]

Hope this helps to find the cause of the invalid parsing.

Greetings

Philipp

I too have issues with this build,

my scenario: Telfort (dutch), vlan4 tv, vlan35 Internet.

The igmpproxy supplied with the 2.3 dailybuild 30dec. stops streaming after ~4minutes.

The version `igmpproxy_20160905_1818.zip` doesn't even shows tv-pictures, and shows (imho) the wrong routes:

--the configfile (generated by pfsense webgui--

```
##-----  
## Enable Quickleave mode (Sends Leave instantly)  
##-----  
quickleave  
phyint sk2 downstream ratelimit 0 threshold 1  
altnet 192.168.89.0/24  
  
phyint sk0_vlan4 upstream ratelimit 0 threshold 1  
altnet 213.75.0.0/16  
altnet 10.142.64.0/18  
  
phyint bridge0 disabled  
phyint sk1 disabled  
phyint sk3 disabled  
phyint sk0_vlan34 disabled
```

--the zipfile version:

```
[2.3.3-DEVELOPMENT][root@pfSense.gijzen.lan]/tmp/bsd10.3_i386/bsd10.3_i386: ./igmpproxy -dvvvv /tmp/igmpproxy.  
conf  
Searching for config file at '/tmp/igmpproxy.conf'  
Config: Quick leave mode enabled.  
Config: Got a phyint token.  
Config: IF: Config for interface sk2.  
Config: IF: Got downstream token.  
Config: IF: Got ratelimit token '0'.  
Config: IF: Got threshold token '1'.  
Config: IF: Got altnet token 192.168.89.0/24.  
Config: IF: Altnet: Parsed altnet to 0.168.89/24.  
IF name : sk2  
Next ptr : 0
```

```

Ratelimit : 0
Threshold : 1
State : 2
Allowednet ptr : 28815030
Config: Got a phyint token.
Config: IF: Config for interface sk0_vlan4.
Config: IF: Got upstream token.
Config: IF: Got ratelimit token '0'.
Config: IF: Got threshold token '1'.
Config: IF: Got altnet token 213.75.0.0/16.
Config: IF: Altnet: Parsed altnet to 0.0/16.
Config: IF: Got altnet token 10.142.64.0/18.
Config: IF: Altnet: Parsed altnet to 0.128.64/18.
IF name : sk0_vlan4
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 1
Allowednet ptr : 28815050
Config: Got a phyint token.
Config: IF: Config for interface bridge0.
Config: IF: Got disabled token.
IF name : bridge0
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface sk1.
Config: IF: Got disabled token.
IF name : sk1
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface sk3.
Config: IF: Got disabled token.
IF name : sk3
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface sk0_vlan34.
Config: IF: Got disabled token.
IF name : sk0_vlan34
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
buildIfVc: Starting...
buildIfVc: Interface is not AF_INET, skipping sk0 (family 18)
buildIfVc: Interface is not AF_INET, skipping sk0 (family 28)
buildIfVc: Interface is not AF_INET, skipping sk1 (family 18)
buildIfVc: Interface is not AF_INET, skipping sk1 (family 28)
buildIfVc: Interface sk1 Addr: 192.168.88.1, Flags: 0xffff8943, Network: 192.168.88/24
buildIfVc: Interface is not AF_INET, skipping sk2 (family 18)
buildIfVc: Interface is not AF_INET, skipping sk2 (family 28)
buildIfVc: Interface sk2 Addr: 192.168.89.1, Flags: 0xffff8843, Network: 192.168.89/24
buildIfVc: Interface is not AF_INET, skipping sk3 (family 18)
buildIfVc: Interface is not AF_INET, skipping sk3 (family 28)
buildIfVc: Interface is not AF_INET, skipping pflog0 (family 18)
buildIfVc: Interface is not AF_INET, skipping pfsync0 (family 18)
buildIfVc: Interface is not AF_INET, skipping enc0 (family 18)
buildIfVc: Interface is not AF_INET, skipping lo0 (family 18)
buildIfVc: Interface lo0 Addr: 127.0.0.1, Flags: 0xffff8049, Network: 127/8
buildIfVc: Interface is not AF_INET, skipping lo0 (family 28)
buildIfVc: Interface is not AF_INET, skipping lo0 (family 28)
buildIfVc: Interface is not AF_INET, skipping sk0_vlan34 (family 18)
buildIfVc: Interface is not AF_INET, skipping sk0_vlan34 (family 28)

```

```

buildIfVc: Interface is not AF_INET, skipping sk0_vlan4 (family 18)
buildIfVc: Interface is not AF_INET, skipping sk0_vlan4 (family 28)
buildIfVc: Interface sk0_vlan4 Addr: 10.235.177.53, Flags: 0xffff8843, Network: 10.235.176/20
buildIfVc: Interface is not AF_INET, skipping bridge0 (family 18)
buildIfVc: Interface bridge0 Addr: 145.129.172.107, Flags: 0xffff8843, Network: 145.129.172/24
buildIfVc: Interface is not AF_INET, skipping bridgel (family 18)
Found config for sk1
Found config for sk2
Found config for sk0_vlan4
Found config for bridge0
adding VIF, Ix 0 Fl 0x0 IP 0x0159a8c0 sk2, Threshold: 1, Ratelimit: 0
    Network for [sk2] : 192.168.89/24
    Network for [sk2] : 0.168.89/24
adding VIF, Ix 1 Fl 0x0 IP 0x35b1eb0a sk0_vlan4, Threshold: 1, Ratelimit: 0
    Network for [sk0_vlan4] : 10.235.176/20
    Network for [sk0_vlan4] : 0.0/16
    Network for [sk0_vlan4] : 0.128.64/18
Got 262144 byte buffer size in 0 iterations
Joining all-routers group 224.0.0.2 on vif 192.168.89.1
joinMcGroup: 224.0.0.2 on sk2
SENT Membership query from 192.168.89.1 to 224.0.0.1
Sent membership query from 192.168.89.1 to 224.0.0.1. Delay: 10
Created timeout 1 (#0) - delay 10 secs
(Id:1, Time:10)
Created timeout 2 (#1) - delay 21 secs
(Id:1, Time:10)
(Id:2, Time:21)
RECV Membership query from 192.168.89.1 to 224.0.0.1
RECV V2 member report from 192.168.89.10 to 224.0.252.132
Should insert group 224.0.252.132 (from: 192.168.89.10) to route table. Vif Ix : 0
No existing route for 224.0.252.132. Create new.
No routes in table. Insert at beginning.
Inserted route table entry for 224.0.252.132 on VIF #0
Joining group 224.0.252.132 upstream on IF address 10.235.177.53
joinMcGroup: 224.0.252.132 on sk0_vlan4

Current routing table (Insert Route):
-----
#0: Src: 0.0.0.0, Dst: 224.0.252.132, Age:2, St: I, OutVifs: 0x00000001
-----
The source address 213.75.167.18 for group 224.0.252.132, is not in any valid net for upstream VIF.
The source address 213.75.167.18 for group 224.0.252.132, is not in any valid net for upstream VIF.
^C
select() failure; Errno(4): Interrupted system call
Got a interrupt signal. Exiting.
clean handler called
Removing route entry for 224.0.252.132
Route is not active. No kernel updates done.
Leaving group 224.0.252.132 upstream on IF address 10.235.177.53
leaveMcGroup: 224.0.252.132 on sk0_vlan4
All routes removed. Routing table is empty.
Shutdown complete....
[2.3.3-DEVELOPMENT][root@pfSense.gijsen.lan]/tmp/bsd10.3_i386/bsd10.3_i386:

```

--the 2.3.3 igmpproxy version dailybuild 30dec --

```

[2.3.3-DEVELOPMENT][root@pfSense.gijsen.lan]/tmp/bsd10.3_i386/bsd10.3_i386: igmpproxy -dvvvv /tmp/igmpproxy.conf
Searching for config file at '/tmp/igmpproxy.conf'
Config: Quick leave mode enabled.
Config: Got a phyint token.
Config: IF: Config for interface sk2.
Config: IF: Got downstream token.
Config: IF: Got ratelimit token '0'.
Config: IF: Got threshold token '1'.
Config: IF: Got altnet token 192.168.89.0/24.
Config: IF: Altnet: Parsed altnet to 192.168.89/24.
IF name : sk2
Next ptr : 0
Ratelimit : 0
Threshold : 1

```

```

State : 2
Allowednet ptr : 28815030
Config: Got a phyint token.
Config: IF: Config for interface sk0_vlan4.
Config: IF: Got upstream token.
Config: IF: Got ratelimit token '0'.
Config: IF: Got threshold token '1'.
Config: IF: Got altnet token 213.75.0.0/16.
Config: IF: Altnet: Parsed altnet to 213.75/16.
Config: IF: Got altnet token 10.142.64.0/18.
Config: IF: Altnet: Parsed altnet to 10.142.64/18.
IF name : sk0_vlan4
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 1
Allowednet ptr : 28815050
Config: Got a phyint token.
Config: IF: Config for interface bridge0.
Config: IF: Got disabled token.
IF name : bridge0
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface sk1.
Config: IF: Got disabled token.
IF name : sk1
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface sk3.
Config: IF: Got disabled token.
IF name : sk3
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface sk0_vlan34.
Config: IF: Got disabled token.
IF name : sk0_vlan34
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface sk0_vlan34.
Config: IF: Got disabled token.
IF name : sk0_vlan34
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
buildIfVc: Interface sk1 Addr: 192.168.88.1, Flags: 0xffff8943, Network: 192.168.88/24
buildIfVc: Interface sk2 Addr: 192.168.89.1, Flags: 0xffff8843, Network: 192.168.89/24
buildIfVc: Interface lo0 Addr: 127.0.0.1, Flags: 0xffff8049, Network: 127/8
buildIfVc: Interface sk0_vlan4 Addr: 10.235.177.53, Flags: 0xffff8843, Network: 10.235.176/20
Found config for sk1
Found config for sk2
Found config for sk0_vlan4
adding VIF, Ix 0 Fl 0x0 IP 0x0158a8c0 sk1, Threshold: 1, Ratelimit: 0
    Network for [sk1] : 192.168.88/24
adding VIF, Ix 1 Fl 0x0 IP 0x0159a8c0 sk2, Threshold: 1, Ratelimit: 0
    Network for [sk2] : 192.168.89/24
    Network for [sk2] : 192.168.89/24
adding VIF, Ix 2 Fl 0x0 IP 0x35b1eb0a sk0_vlan4, Threshold: 1, Ratelimit: 0
    Network for [sk0_vlan4] : 10.235.176/20
    Network for [sk0_vlan4] : 213.75/16
    Network for [sk0_vlan4] : 10.142.64/18
Got 262144 byte buffer size in 0 iterations
Joining all-routers group 224.0.0.2 on vif 192.168.89.1
joinMcGroup: 224.0.0.2 on sk2
SENT Membership query from 192.168.89.1 to 224.0.0.1
Sent membership query from 192.168.89.1 to 224.0.0.1. Delay: 10
Created timeout 1 (#0) - delay 10 secs

```

```
(Id:1, Time:10)
Created timeout 2 (#1) - delay 21 secs
(Id:1, Time:10)
(Id:2, Time:21)
RECV V2 member report from 192.168.89.1 to 224.0.0.2
The IGMP message was from myself. Ignoring.
RECV Membership query from 192.168.89.1 to 224.0.0.1
RECV V2 member report from 192.168.89.10 to 239.255.255.250
Should insert group 239.255.255.250 (from: 192.168.89.10) to route table. Vif Ix : 1
No existing route for 239.255.255.250. Create new.
No routes in table. Insert at beginning.
Inserted route table entry for 239.255.255.250 on VIF #1
Joining group 239.255.255.250 upstream on IF address 10.235.177.53
joinMcGroup: 239.255.255.250 on sk0_vlan4
```

Current routing table (Insert Route):

```
-----
#0: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002
-----
```

```
RECV V2 member report from 192.168.89.1 to 224.0.0.2
The IGMP message was from myself. Ignoring.
RECV V2 member report from 192.168.89.10 to 224.0.252.132
Should insert group 224.0.252.132 (from: 192.168.89.10) to route table. Vif Ix : 1
No existing route for 224.0.252.132. Create new.
Found existing routes. Find insert location.
Inserting at beginning, before route 239.255.255.250
Inserted route table entry for 224.0.252.132 on VIF #1
Joining group 224.0.252.132 upstream on IF address 10.235.177.53
joinMcGroup: 224.0.252.132 on sk0_vlan4
```

Current routing table (Insert Route):

```
-----
#0: Src: 0.0.0.0, Dst: 224.0.252.132, Age:2, St: I, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002
-----
```

```
Route activate request from 213.75.167.18 to 224.0.252.132
Vif bits : 0x00000002
Setting TTL for Vif 1 to 1
Adding MFC: 213.75.167.18 -> 224.0.252.132, InpVif: 2
```

Current routing table (Activate Route):

```
-----
#0: Src: 213.75.167.18, Dst: 224.0.252.132, Age:2, St: A, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002
-----
```

```
^Cselect() failure; Errno(4): Interrupted system call
Got a interrupt signal. Exiting.
clean handler called
Removing route entry for 224.0.252.132
Vif bits : 0x00000002
Setting TTL for Vif 1 to 1
Removing MFC: 213.75.167.18 -> 224.0.252.132, InpVif: 2
Leaving group 224.0.252.132 upstream on IF address 10.235.177.53
leaveMcGroup: 224.0.252.132 on sk0_vlan4
Removing route entry for 239.255.255.250
Route is not active. No kernel updates done.
Leaving group 239.255.255.250 upstream on IF address 10.235.177.53
leaveMcGroup: 239.255.255.250 on sk0_vlan4
All routes removed. Routing table is empty.
```

btw: I use the bridge on my wan-device to get different MACS when obtaining dhcp-leases, my provider won't push multiple ip's (on diff. vlans) with same MAC.

#103 - 12/31/2016 07:25 AM - Vincent Gijzen

edit: 31-12-2016

I've established a working setup, using the develop version, already on the box.
When I disable the filter bogus and filter private ip's, my iptv works consistently ;)

#104 - 01/02/2017 11:39 AM - Luiz Souza

Ooops. Sorry for the breakage.

Fixed in the latest version.

Thanks for the report.

#105 - 01/02/2017 01:41 PM - Renato Botelho

- Assignee changed from Renato Botelho to Luiz Souza

#106 - 01/03/2017 04:44 PM - Philipp Haefelfinger

No problem, sh** happens ;-)

I updated my box today to version 2.4.0.b.20170103.0147.
Checked igmpproxy for new binaries and it looks good.
A short test with the tv-box also looks good and the proxy seems to work. Switching between channels works as well.

Thanks for your good work :-D

#107 - 01/04/2017 05:51 AM - Renato Botelho

- Status changed from Feedback to Resolved

#108 - 01/04/2017 09:00 AM - Lars Veldcholte

Luiz Otavio O Souza wrote:

Ooops. Sorry for the breakage.

Fixed in the latest version.

Thanks for the report.

Where can I download the latest version of IGMPProxy? Or is the only way to upgrade my pfSense to the dev version?

#109 - 01/04/2017 07:34 PM - Alexandre Paradis

Is the change also available to 2.3.3 branch ?

#110 - 01/19/2017 01:19 PM - Lars Veldcholte

Alexandre Paradis wrote:

Is the change also available to 2.3.3 branch ?

I have the same question. Where can I find the latest working version of IGMPProxy? I'm on 2.3.2-RELEASE-p1. Can I download binaries or source (I can compile it myself if needed) of IGMPProxy somewhere, or is the latest version also in the daily snapshots of 2.3.3?

I'm still having this issue where the stream cuts out after a few minutes. It might not be related to IGMPProxy at all but I'd like to test the new version.

#111 - 01/19/2017 01:32 PM - Jeremy Lewis

I have the same issue to IPTV as Alexandre, and have the same version. Is there anyway to update the patch fix for IGMP?

#112 - 01/20/2017 04:20 PM - Greg Myran

Where can I download the latest version of IGMPProxy? Or is the only way to upgrade my pfSense to the dev version?

Links to download the files are at the top of this forum. They have been added to the original bug report.

look for : "igmpproxy_20160905_1818.zip - New version based on the reviewed patch (353 KB) Jorge M. Oliveira, 09/05/2016 12:26 PM"

You need to replace '/usr/local/sbin/igmpproxy' with the appropriate version (bsd10.3_amd64 or bsd10.3_i386) from downloaded zip file *** at your own risk ***

Make sure to backup the original in case something goes wrong and you need to stop the igmpproxy service, replace the file, then restart igmpproxy.

#113 - 01/20/2017 04:28 PM - Lars Veldcholte

Greg Myran wrote:

Where can I download the latest version of IGMPProxy? Or is the only way to upgrade my pfSense to the dev version?

Links to download the files are at the top of this forum. They have been added to the original bug report.

look for : "igmpproxy_20160905_1818.zip - New version based on the reviewed patch (353 KB) Jorge M. Oliveira, 09/05/2016 12:26 PM"

You need to replace '/usr/local/sbin/igmpproxy' with the appropriate version (bsd10.3_amd64 or bsd10.3_i386) from downloaded zip file * **at your own risk** *

Make sure to backup the original in case something goes wrong and you need to stop the igmpproxy service, replace the file, then restart igmpproxy.

That version is from September 2016; I was referring to the fix Luiz Otavio O Souza mentioned a few weeks ago.

I did try that version from the zip though; I have the same problem with both that version and the version from <https://github.com/ViToni/igmpproxy/tree/logging> that the stream cuts out after a while. The route is still in the routing table when that happens though so I have no idea why it happens.

#114 - 01/21/2017 04:15 AM - Harald Gutmann

You can find the sources to build the FreeBSD/pfSense package at: <https://github.com/pfsense/FreeBSD-ports/tree/devel/net/igmpproxy>

The changes related to igmpproxy & this thread are the following commits:

<https://github.com/pfsense/FreeBSD-ports/commit/3cc4c553b9f1697cf1a174c019781abac2254faf>

<https://github.com/pfsense/FreeBSD-ports/commit/009d60cab48f85b0e75a6a35bd7805ece19e6806>

#115 - 01/27/2017 08:58 PM - Andy Shulman

- File *igmpproxy_amd64_2016-01-27.zip* added

I downloaded the 2.3.3 nightly but this fix doesn't appear to be included there, so I built a copy of the proxy from source. It's working well on my 2.3.2-p1 system. I've uploaded a copy of it with today's date. It's only for amd64 (because that's what I use), but I hope some people find this helpful.

EDIT: Sigh, I marked it as 2016-01-27. I clearly haven't figured out what year it is yet.

#116 - 02/15/2017 09:18 AM - Jorge M. Oliveira

Hello,

Okay guys I recognize that one of my latest patches has messed up the behavior earlier. The version that loos-br committed with 009d60cab48f85b0e75a6a35bd7805ece19e6806 is the exact source same as the one on igmpproxy_all.zip so please disregard the newer igmpproxy_20160905_1818.zip as that one is bugged.

Thanks everyone for their feedback and work.

#117 - 02/15/2017 08:40 PM - Kill Bill

I must be extremely dense. Why's this marked as resolved?! Where's a working version available?

#118 - 02/16/2017 10:47 AM - Luiz Souza

It seems to be fixed only on devel branch (2.4 only): <https://github.com/pfsense/FreeBSD-ports/commits/devel/net/igmpproxy/files>

#119 - 02/22/2017 02:12 AM - Joao Dinis Neves

Do you guys know if this is fixed on version 2.3.3.

Thanks.

#120 - 02/22/2017 02:57 AM - Kill Bill

Joao Dinis Neves wrote:

Do you guys know if this is fixed on version 2.3.3.

Thanks.

Read the comment right above yours.

#121 - 03/01/2017 08:46 AM - Rai Wol

Can someone confirm its working in 2.4?

Doesn't stop after 3-4 min?

#122 - 03/01/2017 09:28 AM - Matthias Lange

Hi,

I can confirm that the bug fix is working (at least for me: IPTV in Germany, no BNG)

Runs since two weeks without any interrupts

#123 - 03/01/2017 11:22 AM - Rai Wol

Just upgraded to 2.4.

Interfaces are recognized correctly. So actual bug is fixed.

But multicast stream still cuts out after 3-4 min. Last log entries shows that the membership report is received correctly:

```
Mar 1 17:55:56 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:56:16 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:56:36 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:56:56 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:57:16 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:57:36 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:57:46 pfSense igmpproxy: RECV V2 member report from 192.168.5.101 to 239.255.255.250
Mar 1 17:57:47 pfSense igmpproxy: RECV V2 member report from 192.168.5.101 to 224.0.251.134
```

#124 - 03/02/2017 12:06 AM - Philipp Haefelfinger

I see the same messages on my system but my streams do not cut out after any time. They are really stable and channelswitching is working without a problem.

#125 - 03/02/2017 11:29 AM - Rai Wol

Rai Wol wrote:

Just upgraded to 2.4.

Interfaces are recognized correctly. So actual bug is fixed.

But multicast stream still cuts out after 3-4 min. Last log entries shows that the membership report is received correctly:

```
Mar 1 17:55:56 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:56:16 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:56:36 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:56:56 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:57:16 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:57:36 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:57:46 pfSense igmpproxy: RECV V2 member report from 192.168.5.101 to 239.255.255.250
Mar 1 17:57:47 pfSense igmpproxy: RECV V2 member report from 192.168.5.101 to 224.0.251.134
```

Nevermind. Problem was the firewall. Advanced IP options was not enabled on upstream interface. After enabled IPTV is working!

#126 - 03/02/2017 05:34 PM - Lars Veldcholte

Rai Wol wrote:

Rai Wol wrote:

Just upgraded to 2.4.

Interfaces are recognized correctly. So actual bug is fixed.

But multicast stream still cuts out after 3-4 min. Last log entries shows that the membership report is received correctly:

```
Mar 1 17:55:56 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:56:16 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:56:36 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:56:56 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:57:16 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
Mar 1 17:57:36 pfSense igmpproxy: The source address 192.168.5.101 for group 239.255.255.250, is not in any valid net for upstream VIF.
```

Mar 1 17:57:46 pfSense igmpproxy: RECV V2 member report from 192.168.5.101 to 239.255.255.250
Mar 1 17:57:47 pfSense igmpproxy: RECV V2 member report from 192.168.5.101 to 224.0.251.134

Nevermind. Problem was the firewall. Advanced IP options was not enabled on upstream interface. After enabled IPTV is working!

I think I have the same problem, except that it takes somewhat longer for the stream to cut out (around 15 min) in my case. Could you elaborate on what fixed your problem? What firewall rules do you have? I have an "allow multicast" rule on my upstream interface, which allows UDP traffic to 224.0.0.0/4, with "allow IP options" checked. I have a "allow all" rule on my downstream interface, also with "allow IP options" checked. Do I need anything else?

#127 - 05/18/2017 04:53 PM - Diogo Quintela

Rai Wol wrote:

Can someone confirm its working in 2.4?

Doesn't stop after 3-4 min?

igmpproxy_all.zip made it working to my pfsense 2.3.4 - working in PT's MEO IPTV.

Is there any change this fix get merged on 2.3 branch ? Does a PR be made explicit for this ?
I have a interest in that, because I don't know if my box will support 2.4 (freebsd 11) - tried install freebsd there and freezed on install because of weird acpi bios problems. 2.3's freebsd 10 complained but made it work.

Regards
Diogo Quintela

#128 - 06/22/2017 04:18 PM - Mr B

Hi,
This still isn't working for me on 2.4 - 2.4.0.b.20170622.0342 - keep getting the cut off after 4 minutes.

In the log i cant see anything wrong apart from the routes aging - but the connection still cuts out...

@[2.4.0-BETA][admin@pfsense.homenet.local]/tmp: cat igmpproxy.conf

##-----

1. Enable Quickleave mode (Sends Leave instantly)

##-----

```
quickleave
phyint vmx2_vlan20 downstream ratelimit 0 threshold 1
altnet 192.168.20.1/30
```

```
phyint vmx1 upstream ratelimit 0 threshold 1
altnet 224.0.0.0/4
altnet 109.159.247.0/24
```

```
phyint pppoe1 disabled
phyint vmx2 disabled
phyint vmx3 disabled
phyint vmx0 disabled
phyint ovpcnc1 disabled
phyint vmx2_vlan30 disabled
```

[2.4.0-BETA][admin@pfsense.homenet.local]/tmp: igmpproxy -d -vvvv /tmp/igmpproxy.conf

Searching for config file at '/tmp/igmpproxy.conf'

Config: Quick leave mode enabled.

Config: Got a phyint token.

Config: IF: Config for interface vmx2_vlan20.

Config: IF: Got downstream token.

Config: IF: Got ratelimit token '0'.

Config: IF: Got threshold token '1'.

Config: IF: Got altnet token 192.168.20.1/30.

Config: IF: Altnet: Parsed altnet to 192.168.20.1/30.

IF name : vmx2_vlan20

Next ptr : 0

Ratelimit : 0

Threshold : 1

State : 2

Allowednet ptr : e24010

Config: Got a phyint token.

Config: IF: Config for interface vmx1.

Config: IF: Got upstream token.

Config: IF: Got ratelimit token '0'.

Config: IF: Got threshold token '1'.

Config: IF: Got altnet token 224.0.0.0/4.

Config: IF: Altnet: Parsed altnet to 224/4.

Config: IF: Got altnet token 109.159.247.0/24.

Config: IF: Altnet: Parsed altnet to 109.159.247/24.

IF name : vmx1

Next ptr : 0

Ratelimit : 0

Threshold : 1

State : 1

Allowednet ptr : e24020

Config: Got a phyint token.

Config: IF: Config for interface pppoe1.

Config: IF: Got disabled token.

IF name : pppoe1

Next ptr : 0

Ratelimit : 0

Threshold : 1

State : 0

Allowednet ptr : 0

Config: Got a phyint token.

Config: IF: Config for interface vmx2.

Config: IF: Got disabled token.

IF name : vmx2

Next ptr : 0

Ratelimit : 0

Threshold : 1

State : 0

Allowednet ptr : 0

Config: Got a phyint token.

Config: IF: Config for interface vmx3.

Config: IF: Got disabled token.

IF name : vmx3

Next ptr : 0

Ratelimit : 0

Threshold : 1

State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface vmx0.
Config: IF: Got disabled token.
IF name : vmx0
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface ovpc1.
Config: IF: Got disabled token.
IF name : ovpc1
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
Config: Got a phyint token.
Config: IF: Config for interface vmx2_vlan30.
Config: IF: Got disabled token.
IF name : vmx2_vlan30
Next ptr : 0
Ratelimit : 0
Threshold : 1
State : 0
Allowednet ptr : 0
buildfVc: Starting...
buildfVc: Interface is not AF_INET, skipping vmx0 (family 18)
buildfVc: Interface is not AF_INET, skipping vmx0 (family 28)
buildfVc: Interface vmx0 Addr: 192.168.1.1, Flags: 0xffff8943, Network: 192.168.1/24
buildfVc: Interface is not AF_INET, skipping vmx1 (family 18)
buildfVc: Interface is not AF_INET, skipping vmx1 (family 28)
buildfVc: Interface vmx1 Addr: 10.20.30.40, Flags: 0xffff8843, Network: 10.20.30.40/32
buildfVc: Interface is not AF_INET, skipping vmx2 (family 18)
buildfVc: Interface is not AF_INET, skipping vmx2 (family 28)
buildfVc: Interface vmx2 Addr: 192.168.254.1, Flags: 0xffff8943, Network: 192.168.254/24
buildfVc: Interface is not AF_INET, skipping vmx3 (family 18)
buildfVc: Interface is not AF_INET, skipping vmx3 (family 28)
buildfVc: Interface vmx3 Addr: 192.168.0.1, Flags: 0xffff8943, Network: 192.168.0/24
buildfVc: Interface is not AF_INET, skipping enc0 (family 18)
buildfVc: Interface is not AF_INET, skipping lo0 (family 18)
buildfVc: Interface is not AF_INET, skipping lo0 (family 28)
buildfVc: Interface is not AF_INET, skipping lo0 (family 28)
buildfVc: Interface lo0 Addr: 127.0.0.1, Flags: 0xffff8049, Network: 127/8
buildfVc: Interface is not AF_INET, skipping pfsync0 (family 18)
buildfVc: Interface is not AF_INET, skipping pflog0 (family 18)
buildfVc: Interface is not AF_INET, skipping vmx2_vlan20 (family 18)
buildfVc: Interface is not AF_INET, skipping vmx2_vlan20 (family 28)
buildfVc: Interface vmx2_vlan20 Addr: 192.168.20.1, Flags: 0xffff8843, Network: 192.168.20.0/30
buildfVc: Interface is not AF_INET, skipping vmx2_vlan30 (family 18)
buildfVc: Interface is not AF_INET, skipping vmx2_vlan30 (family 28)
buildfVc: Interface vmx2_vlan30 Addr: 192.168.30.1, Flags: 0xffff8843, Network: 192.168.30/24
buildfVc: Interface is not AF_INET, skipping pppoe1 (family 18)
buildfVc: Interface is not AF_INET, skipping pppoe1 (family 28)
buildfVc: Interface pppoe1 Addr: xxxxxxxxxxxx, Flags: 0xffff89d1, Network: xxxxxxxxxxxx
buildfVc: Interface is not AF_INET, skipping ovps2 (family 18)
buildfVc: Interface is not AF_INET, skipping ovps2 (family 28)
buildfVc: Interface ovps2 Addr: 10.8.5.1, Flags: 0xffff8051, Network: 10.8.5/24
buildfVc: Interface is not AF_INET, skipping ovpc1 (family 18)
buildfVc: Interface is not AF_INET, skipping ovpc1 (family 28)
buildfVc: Interface ovpc1 Addr: 10.4.2.227, Flags: 0xffff8051, Network: 10.4/16
Found config for vmx0
Found config for vmx1
Found config for vmx2
Found config for vmx3
Found config for vmx2_vlan20
Found config for vmx2_vlan30
Found config for pppoe1
Found config for ovpc1
adding VIF, Ix 0 FI 0x0 IP 0x281e140a vmx1, Threshold: 1, Ratelimit: 0
Network for [vmx1] : 10.20.30.40/32
Network for [vmx1] : 224/4

Network for [vmx1] : 109.159.247/24
adding VIF, Ix 1 FI 0x0 IP 0x0114a8c0 vmx2_vlan20, Threshold: 1, Ratelimit: 0
Network for [vmx2_vlan20] : 192.168.20.0/30
Network for [vmx2_vlan20] : 192.168.20.1/30
adding VIF, Ix 2 FI 0x0 IP 0x0105080a ovpsns2, Threshold: 1, Ratelimit: 0
Network for [ovpsns2] : 10.8.5/24
Got 262144 byte buffer size in 0 iterations
Joining all-routers group 224.0.0.2 on vif 192.168.20.1
joinMcGroup: 224.0.0.2 on vmx2_vlan20
Joining all-routers group 224.0.0.2 on vif 10.8.5.1
joinMcGroup: 224.0.0.2 on ovpsns2
SENT Membership query from 192.168.20.1 to 224.0.0.1
Sent membership query from 192.168.20.1 to 224.0.0.1. Delay: 10
SENT Membership query from 10.8.5.1 to 224.0.0.1
Sent membership query from 10.8.5.1 to 224.0.0.1. Delay: 10
Created timeout 1 (#0) - delay 10 secs
(Id:1, Time:10)
Created timeout 2 (#1) - delay 21 secs
(Id:1, Time:10)
(Id:2, Time:21)
RECV Membership query from 192.168.20.1 to 224.0.0.1
RECV Membership query from 10.8.5.1 to 224.0.0.1
RECV V2 member report from 192.168.20.1 to 224.0.0.2
The IGMP message was from myself. Ignoring.
RECV V2 member report from 192.168.20.2 to 239.255.255.250
Should insert group 239.255.255.250 (from: 192.168.20.2) to route table. Vif Ix : 1
No existing route for 239.255.255.250. Create new.
No routes in table. Insert at beginning.
Inserted route table entry for 239.255.255.250 on VIF #1
Joining group 239.255.255.250 upstream on IF address 10.20.30.40
joinMcGroup: 239.255.255.250 on vmx1

Current routing table (Insert Route):

#0: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002

RECV V2 member report from 192.168.20.2 to 234.81.130.36
Should insert group 234.81.130.36 (from: 192.168.20.2) to route table. Vif Ix : 1
No existing route for 234.81.130.36. Create new.
Found existing routes. Find insert location.
Inserting at beginning, before route 239.255.255.250
Inserted route table entry for 234.81.130.36 on VIF #1
Joining group 234.81.130.36 upstream on IF address 10.20.30.40
joinMcGroup: 234.81.130.36 on vmx1

Current routing table (Insert Route):

#0: Src: 0.0.0.0, Dst: 234.81.130.36, Age:2, St: I, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002

Route activate request from 109.159.247.1 to 234.81.130.36
Vif bits : 0x00000002
Setting TTL for Vif 1 to 1
Adding MFC: 109.159.247.1 -> 234.81.130.36, InpVif: 0

Current routing table (Activate Route):

#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:2, St: A, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002

RECV V2 member report from 192.168.20.2 to 224.0.0.251
Should insert group 224.0.0.251 (from: 192.168.20.2) to route table. Vif Ix : 1
No existing route for 224.0.0.251. Create new.
Found existing routes. Find insert location.
Inserting after route 239.255.255.250
Inserted route table entry for 224.0.0.251 on VIF #1
Joining group 224.0.0.251 upstream on IF address 10.20.30.40
joinMcGroup: 224.0.0.251 on vmx1

Current routing table (Insert Route):

#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:2, St: A, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002
#2: Src: 0.0.0.0, Dst: 224.0.0.251, Age:2, St: I, OutVifs: 0x00000002

About to call timeout 1 (#0)
Aging routes in table.

Current routing table (Age active routes):

#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:1, St: A, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:1, St: I, OutVifs: 0x00000002
#2: Src: 0.0.0.0, Dst: 224.0.0.251, Age:1, St: I, OutVifs: 0x00000002

About to call timeout 2 (#0)
SENT Membership query from 192.168.20.1 to 224.0.0.1
Sent membership query from 192.168.20.1 to 224.0.0.1. Delay: 10
SENT Membership query from 10.8.5.1 to 224.0.0.1
Sent membership query from 10.8.5.1 to 224.0.0.1. Delay: 10
Created timeout 3 (#0) - delay 10 secs
(Id:3, Time:10)
Created timeout 4 (#1) - delay 21 secs
(Id:3, Time:10)
(Id:4, Time:21)
RECV Membership query from 192.168.20.1 to 224.0.0.1
RECV Membership query from 10.8.5.1 to 224.0.0.1
RECV V2 member report from 192.168.20.2 to 224.0.0.251
Should insert group 224.0.0.251 (from: 192.168.20.2) to route table. Vif Ix : 1
Updated route entry for 224.0.0.251 on VIF #1

Current routing table (Insert Route):

#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:1, St: A, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:1, St: I, OutVifs: 0x00000002
#2: Src: 0.0.0.0, Dst: 224.0.0.251, Age:1, St: I, OutVifs: 0x00000002

RECV V2 member report from 192.168.20.2 to 239.255.255.250
Should insert group 239.255.255.250 (from: 192.168.20.2) to route table. Vif Ix : 1
Updated route entry for 239.255.255.250 on VIF #1

Current routing table (Insert Route):

#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:1, St: A, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:1, St: I, OutVifs: 0x00000002
#2: Src: 0.0.0.0, Dst: 224.0.0.251, Age:1, St: I, OutVifs: 0x00000002

RECV V2 member report from 192.168.20.1 to 224.0.0.2
The IGMP message was from myself. Ignoring.
RECV V2 member report from 192.168.20.2 to 234.81.130.36
Should insert group 234.81.130.36 (from: 192.168.20.2) to route table. Vif Ix : 1
Updated route entry for 234.81.130.36 on VIF #1
Vif bits : 0x00000002
Setting TTL for Vif 1 to 1
Adding MFC: 109.159.247.1 -> 234.81.130.36, InpVif: 0

Current routing table (Insert Route):

#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:1, St: A, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:1, St: I, OutVifs: 0x00000002
#2: Src: 0.0.0.0, Dst: 224.0.0.251, Age:1, St: I, OutVifs: 0x00000002

About to call timeout 3 (#0)
Aging routes in table.

Current routing table (Age active routes):

#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:2, St: A, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002
#2: Src: 0.0.0.0, Dst: 224.0.0.251, Age:2, St: I, OutVifs: 0x00000002

About to call timeout 4 (#0)
SENT Membership query from 192.168.20.1 to 224.0.0.1
Sent membership query from 192.168.20.1 to 224.0.0.1. Delay: 10
SENT Membership query from 10.8.5.1 to 224.0.0.1
Sent membership query from 10.8.5.1 to 224.0.0.1. Delay: 10
Created timeout 5 (#0) - delay 10 secs
(Id:5, Time:10)
Created timeout 6 (#1) - delay 115 secs
(Id:5, Time:10)
(Id:6, Time:115)

RECV Membership query from 192.168.20.1 to 224.0.0.1
RECV Membership query from 10.8.5.1 to 224.0.0.1
RECV V2 member report from 192.168.20.2 to 224.0.0.251
Should insert group 224.0.0.251 (from: 192.168.20.2) to route table. Vif lx : 1
Updated route entry for 224.0.0.251 on VIF #1

Current routing table (Insert Route):

#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:2, St: A, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002
#2: Src: 0.0.0.0, Dst: 224.0.0.251, Age:2, St: I, OutVifs: 0x00000002

RECV V2 member report from 192.168.20.2 to 234.81.130.36
Should insert group 234.81.130.36 (from: 192.168.20.2) to route table. Vif lx : 1
Updated route entry for 234.81.130.36 on VIF #1
Vif bits : 0x00000002
Setting TTL for Vif 1 to 1
Adding MFC: 109.159.247.1 -> 234.81.130.36, InpVif: 0

Current routing table (Insert Route):

#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:2, St: A, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002
#2: Src: 0.0.0.0, Dst: 224.0.0.251, Age:2, St: I, OutVifs: 0x00000002

RECV V2 member report from 192.168.20.2 to 239.255.255.250
Should insert group 239.255.255.250 (from: 192.168.20.2) to route table. Vif lx : 1
Updated route entry for 239.255.255.250 on VIF #1

Current routing table (Insert Route):

#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:2, St: A, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002
#2: Src: 0.0.0.0, Dst: 224.0.0.251, Age:2, St: I, OutVifs: 0x00000002

RECV V2 member report from 192.168.20.1 to 224.0.0.2
The IGMP message was from myself. Ignoring.
About to call timeout 5 (#0)
Aging routes in table.

Current routing table (Age active routes):

#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:2, St: A, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002
#2: Src: 0.0.0.0, Dst: 224.0.0.251, Age:2, St: I, OutVifs: 0x00000002

About to call timeout 6 (#0)
SENT Membership query from 192.168.20.1 to 224.0.0.1
Sent membership query from 192.168.20.1 to 224.0.0.1. Delay: 10
SENT Membership query from 10.8.5.1 to 224.0.0.1
Sent membership query from 10.8.5.1 to 224.0.0.1. Delay: 10
Created timeout 7 (#0) - delay 10 secs
(Id:7, Time:10)
Created timeout 8 (#1) - delay 115 secs
(Id:7, Time:10)
(Id:8, Time:115)
RECV Membership query from 192.168.20.1 to 224.0.0.1
RECV Membership query from 10.8.5.1 to 224.0.0.1
RECV V2 member report from 192.168.20.2 to 234.81.130.36
Should insert group 234.81.130.36 (from: 192.168.20.2) to route table. Vif lx : 1
Updated route entry for 234.81.130.36 on VIF #1
Vif bits : 0x00000002
Setting TTL for Vif 1 to 1
Adding MFC: 109.159.247.1 -> 234.81.130.36, InpVif: 0

Current routing table (Insert Route):

#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:2, St: A, OutVifs: 0x00000002
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002
#2: Src: 0.0.0.0, Dst: 224.0.0.251, Age:2, St: I, OutVifs: 0x00000002

RECV V2 member report from 192.168.20.2 to 224.0.0.251
Should insert group 224.0.0.251 (from: 192.168.20.2) to route table. Vif lx : 1
Updated route entry for 224.0.0.251 on VIF #1

Current routing table (Insert Route):

```
-----  
#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:2, St: A, OutVifs: 0x00000002  
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002  
#2: Src: 0.0.0.0, Dst: 224.0.0.251, Age:2, St: I, OutVifs: 0x00000002  
-----
```

```
RECV V2 member report from 192.168.20.1 to 224.0.0.2  
The IGMP message was from myself. Ignoring.  
RECV V2 member report from 192.168.20.2 to 239.255.255.250  
Should insert group 239.255.255.250 (from: 192.168.20.2) to route table. Vif ix : 1  
Updated route entry for 239.255.255.250 on VIF #1
```

Current routing table (Insert Route):

```
-----  
#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:2, St: A, OutVifs: 0x00000002  
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002  
#2: Src: 0.0.0.0, Dst: 224.0.0.251, Age:2, St: I, OutVifs: 0x00000002  
-----
```

```
About to call timeout 7 (#0)  
Aging routes in table.
```

Current routing table (Age active routes):

```
-----  
#0: Src: 109.159.247.1, Dst: 234.81.130.36, Age:2, St: A, OutVifs: 0x00000002  
#1: Src: 0.0.0.0, Dst: 239.255.255.250, Age:2, St: I, OutVifs: 0x00000002  
#2: Src: 0.0.0.0, Dst: 224.0.0.251, Age:2, St: I, OutVifs: 0x00000002  
-----
```

@

Does anyone have any pointers for me here?

#129 - 06/22/2017 06:13 PM - Jeremy Lewis

The way I managed to get it working reliably was to turn off the IGMP snooping on my managed switch, then the timing out issue went away, but to prevent the multicast flooding my network I put the Youview box onto a separate VLAN and it all worked a treat.

Since then I've got the Netgate SG-1000 2.4.0-BETA (arm) but had to enable the switch's snooping (with IGMP Querier active protocol v2) again to work on that for to work? The VLAN is still in place but I could probably move the Youview box into my normal LAN now as the Switch is handling Multicast nicely now with the New SG-1000.

#130 - 06/24/2017 03:03 PM - Mr B

Jeremy Lewis wrote:

The way I managed to get it working reliably was to turn off the IGMP snooping on my managed switch, then the timing out issue went away, but to prevent the multicast flooding my network I put the Youview box onto a separate VLAN and it all worked a treat.

Since then I've got the Netgate SG-1000 2.4.0-BETA (arm) but had to enable the switch's snooping (with IGMP Querier active protocol v2) again to work on that for to work? The VLAN is still in place but I could probably move the Youview box into my normal LAN now as the Switch is handling Multicast nicely now with the New SG-1000.

Sorry - resolved my issue - explained here:-

<https://forum.pfsense.org/index.php?topic=100444.msg729455#msg729455>

#131 - 07/05/2017 07:12 PM - J L

Diogo Quintela wrote:

Rai Wol wrote:

Can someone confirm its working in 2.4?

Doesn't stop after 3-4 min?

igmpoxy_all.zip made it working to my pfsense 2.3.4 - working in PT's MEO IPTV.

Is there any change this fix get merged on 2.3 branch ? Does a PR be made explicit for this ?

No idea about it getting merged, someone should work on that.

igmpoxy_all.zip also works on 2.3.3_1 with TELUS IPTV.

#132 - 08/06/2017 09:24 AM - Mr J

—

#133 - 08/06/2017 02:21 PM - Harald Gutmann

Mr J wrote:

Maybe Luiz Otavio O Souza should just FIX the bug for version 2.40 ????!!!!

Probably its not the best idea to come over and shout out what someone should do. Especially when the fault is most likely to be found in your firewall rules.

Let me guess: You didn't create the required multicast rule which allows to pass the necessary IGMP control messages:

<https://redmine.pfsense.org/issues/6099#note-130>

<https://forum.pfsense.org/index.php?PHPSESSID=iosula521i677qfo2v71rhu6s3&topic=100444.msg729455#msg729455>

#134 - 08/06/2017 03:05 PM - Mr J

—

#135 - 08/07/2017 03:43 AM - Harald Gutmann

The way I use igmpproxy it works properly. There are many posts above from other people who confirmed that IPTV works nicely with the fixed version, when the configuration is done properly.

I suggest to double check the configuration and make sure that the follow-up IGMP messages, which typically appear after 3-5 minutes reach the device which tries to see the stream. Make sure "Advanced Options -> Allow packets with IP options" is set in your rule. This option is hidden in the default view.

This bug is/was related to recognizing upstream interfaces, and the effect was that igmpproxy was not able to handle all interfaces when it appeared. Clearly this is not the case with your setup, since it works for the first minutes, which means that igmpproxy does recognize and handle your interfaces.

#136 - 08/07/2017 06:35 AM - Mr J

—

#137 - 08/07/2017 07:08 AM - Jim Pingle

Please take this discussion to the forum, mailing list, reddit, etc. If it runs for a few minutes then it is absolutely unrelated to this ticket. If, after discussing it in an appropriate place, a new bug is identified, then a new issue can be opened for it with more relevant detail at that time.

Thanks!

#138 - 08/07/2017 07:13 AM - Mr J

igmpproxy does indeed work 100% for UK BT TV, BT Sport 4K (IPTV over BT Infinity FTTC/P) in ver. 2.4.0-BETA (Version 2.4.0.b.20170807.0351)

None of the replacement files at the top of this page are needed.

Thanks to Luiz ;-)

#139 - 08/07/2017 12:51 PM - Victor Toni

Some of these reports seem to miss one very important information: which version of IGMP is used.

IMHO igmpproxy does not work well with IGMPv3 and if you need IPv6/MLD you are completely lost...

To sum it up there are basically two issues with igmpproxy:

- IF lookup uses an old mechanism which does not cope well with lots of IFs (physical/virtual ones, etc.)
- IGMPv3 support is not complete and therefor you might have issues (such as some of the customres of Deutsche Telekom if they are using their newer IPTV product EntertainTV)

#140 - 08/08/2017 03:22 PM - Harald Gutmann

For the records:

Mr J wrote on <https://forum.pfsense.org/index.php?topic=134795.0>

These instructions by james_h DO work (Reply #3): <https://forum.pfsense.org/index.php?topic=74126.0>

Just this ONE EXTRA STEP is needed though:

"For my BT UltraHD youview box i also needed another pass rule on the IPTV Wan interface we create.

This should allow IPV4 IGMP, from a source of any, to destination of network 224.0.0.0/4 (with Advanced Options -> "Allow packets with >IP options to pass. Otherwise they are blocked by default. This is usually only seen with multicast traffic" ticked) otherwise the >picture would drop out after 4 minutes."

For 2.4.0-BETA, igmpproxy IS working correctly - no updates/fixes are needed.

So anyone coming here and having troubles with IPTV and igmpproxy please do consider this setting. And kindly re-check it before posting any message to this thread and related to igmpproxy not recognizing the interface. As already stated above, if it works for 3-5 minutes, you are not facing the above discussed issue.

#141 - 10/24/2017 03:10 PM - Samuel Kadolph

Was this fixed in 2.4.0? I updated to it but my IGMP Proxy service is not working with the same message of:

There must be at least 2 Vif's where one is upstream.

Files

igmpproxy.conf	597 Bytes	04/09/2016	Victor Toni
igmpproxy_stdout.log	3.32 KB	04/09/2016	Victor Toni
router.skeba.de - Status_Dashboard_2016-04-28_16-14-15.png	20.8 KB	04/28/2016	Stefan Heck
Dok2.pdf	339 KB	04/28/2016	Stefan Heck
patch-src__os-freebsd.h	486 Bytes	09/04/2016	Jorge M. Oliveira
patch-src_ifvc.c	3.12 KB	09/04/2016	Jorge M. Oliveira
patch-src_igmpproxy.c	362 Bytes	09/04/2016	Jorge M. Oliveira
igmproxy_all.zip	247 KB	09/05/2016	Jorge M. Oliveira
igmpproxy_20160905_1818.zip	353 KB	09/05/2016	Jorge M. Oliveira
igmpproxy_amd64_2016-01-27.zip	17.3 KB	01/28/2017	Andy Shulman