

pfSense - Bug #4446

IP Alias with CARP VIP parent is not removed from OS on secondary node when deleted

02/19/2015 02:03 AM - Andreas Pflug

Status:	Resolved	Start date:	02/19/2015
Priority:	Normal	Due date:	
Assignee:	Renato Botelho	% Done:	100%
Category:	CARP	Estimated time:	0.00 hour
Target version:	2.2.1	Affected Version:	2.2
Plus Target Version:		Affected Architecture:	
Release Notes:	Default		

Description

I used to have a CARP VIP for any IP address my pf cluster has to handle since that used to be the only way (changed in 2.1 I believe?).

With 2.2, I've been defining a single IP Alias on the WAN's CARP VIP, and had trouble with the backup router flipping its WAN VIP to master, resulting in a dual-master situation and intermittent problems. Redefining the VIP to CARP made the problem disappear immediately.

Associated revisions

Revision 457d9c32 - 03/09/2015 12:52 PM - Ermal Luçi

Fixes #4446 Correct ipalias removal on top of carp during configuration sync.

Revision 8896fe1c - 03/09/2015 12:52 PM - Ermal Luçi

Fixes #4446 Correct ipalias removal on top of carp during configuration sync.

Revision 0e7954b8 - 03/10/2015 07:00 PM - Chris Buechler

Use a consistent variable name here. related to Ticket #4446

Revision 309162b9 - 03/10/2015 07:01 PM - Chris Buechler

Use a consistent variable name here. related to Ticket #4446

Conflicts:

usr/local/www/xmlrpc.php

Revision 04bbb3a5 - 03/11/2015 02:02 PM - Renato Botelho

Revert "Use a consistent variable name here. related to Ticket #4446"

This reverts commit 309162b96d6ad0db09c132cc0aabc890bfe1909f.

Revision c14781e3 - 03/11/2015 02:02 PM - Renato Botelho

Save vip interface and subnet to use to delete old vips from secondary nodes. Fixes #4446

Revision b01d44a3 - 03/11/2015 02:02 PM - Renato Botelho

Revert "Use a consistent variable name here. related to Ticket #4446"

This reverts commit 0e7954b8a333d7ca92f56c86c74e2d9d7457b546.

Save vip interface and subnet to use to delete old vips from secondary nodes. Fixes #4446

History

#1 - 02/19/2015 04:43 PM - Chris Buechler

- Status changed from New to Feedback

this definitely works in general. when you end up in dual master is if it doesn't sync the alias across for some reason, or one I've seen where that IP was manually added on the secondary as an IP alias of an interface it fails to add it as an alias of the CARP, which gives you dual master.

likely this is a misconfiguration along those lines. if it's not, there is some circumstance more specific than described that's triggering it.

#2 - 02/20/2015 01:04 AM - Andreas Pflug

Chris Buechler wrote:

this definitely works in general. when you end up in dual master is if it doesn't sync the alias across for some reason, or one I've seen where that IP was manually added on the secondary as an IP alias of an interface it fails to add it as an alias of the CARP, which gives you dual master.

While trying to solve the problem, I manually removed all VIFs (carp as well as alias) from the secondary's config.xml and rebooted. After that, I re-created the VIFs by triggering a sync from the master, so I made sure there's no backup specific config.

To make it clear: the dual-master situation is on the CARP interface that's used for IP alias; removing the IP alias restored normal master-backup CARP.

What info do you need to investigate "some circumstances more specific" which seem to apply here?

#3 - 02/20/2015 02:34 AM - Chris Buechler

what does the output of ifconfig show on the secondary?

#4 - 02/20/2015 02:58 AM - Andreas Pflug

Ok, I know how to reproduce this.

On the master, I have this config:

```
xn0: flags=8943<UP,BROADCAST,RUNNING,PROMISC,SIMPLEX,MULTICAST> metric 0 mtu 1500
    inet xx.xx.xx.91 netmask 0xffffffff broadcast xx.xx.xx.95
    inet xx.xx.xx.90 netmask 0xffffffff broadcast xx.xx.xx.95 vhid 90
    inet xx.xx.xx.94 netmask 0xffffffff broadcast xx.xx.xx.95 vhid 94
    carp: MASTER vhid 90 advbase 1 advskew 0
    carp: MASTER vhid 94 advbase 1 advskew 0
```

No config is performed on the backup, just syncing over.

```
xn0: flags=8943<UP,BROADCAST,RUNNING,PROMISC,SIMPLEX,MULTICAST> metric 0 mtu 1500
    inet xx.xx.xx.92 netmask 0xffffffff broadcast xx.xx.xx.95
    inet xx.xx.xx.90 netmask 0xffffffff broadcast xx.xx.xx.95 vhid 90
    inet xx.xx.xx.94 netmask 0xffffffff broadcast xx.xx.xx.95 vhid 94
    carp: BACKUP vhid 90 advbase 1 advskew 100
    carp: BACKUP vhid 94 advbase 1 advskew 100
```

Looks as expected.

Now I create an IP alias on vhid94, and apply. Things still work ok.

I delete the IP alias. The page won't show "Apply settings now" so I don't do it; the alias is gone immediately, but now CARP on vhid94 is double-master.

I'm triggering a sync by saving CARP settings, no improvement; the backup will show this:

```
xn0: flags=8943<UP,BROADCAST,RUNNING,PROMISC,SIMPLEX,MULTICAST> metric 0 mtu 1500
    inet xx.xx.xx.92 netmask 0xffffffff broadcast xx.xx.xx.95
    inet xx.xx.xx.90 netmask 0xffffffff broadcast xx.xx.xx.95 vhid 90
    inet xx.xx.xx.94 netmask 0xffffffff broadcast xx.xx.xx.95 vhid 94
    inet xx.xx.xx.93 netmask 0xffffffff broadcast xx.xx.xx.95 vhid 94
    carp: BACKUP vhid 90 advbase 1 advskew 100
    carp: MASTER vhid 94 advbase 1 advskew 100
```

Apparently, the deletion of the ip alias wasn't promoted correctly to the secondary.

#5 - 02/20/2015 08:45 AM - Jim Pingle

- Subject changed from double CARP master if used for IP alias to IP Alias with CARP VIP parent is not removed from OS on secondary node when deleted

- Status changed from Feedback to Confirmed

Updated the description to be more accurate. The actual problem appears to be that deleting an IP Alias VIP with a CARP parent from the primary does not also delete that VIP from the actual interface on the secondary. The GUI entry is gone but if you check ifconfig the entry is still present on the secondary.

Because the secondary has an IP address using the VHID that the master does not, it believes it should be master instead (at least for that VIP).

The the example below, I added 192.168.71.5 as an IP alias VIP using the CARP VIP as the parent, and then removed it.

```
root@primary: ifconfig em1
em1: flags=8943<UP,BROADCAST,RUNNING,PROMISC,SIMPLEX,MULTICAST> metric 0 mtu 1500
  options=9b<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, VLAN_HWCSUM>
  ether 00:0c:29:9b:d3:d5
  inet 192.168.71.2 netmask 0xffffffff broadcast 192.168.71.255
  inet6 fe80::20c:29ff:fe9b:d3d5%em1 prefixlen 64 scopeid 0x2
  inet 192.168.71.1 netmask 0xffffffff broadcast 192.168.71.255 vhid 72
  nd6 options=21<PERFORMNUD,AUTO_LINKLOCAL>
  media: Ethernet autoselect (1000baseT <full-duplex>)
  status: active
  carp: MASTER vhid 72 advbase 1 advskew 0

root@secondary: ifconfig em1
em1: flags=8943<UP,BROADCAST,RUNNING,PROMISC,SIMPLEX,MULTICAST> metric 0 mtu 1500
  options=9b<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, VLAN_HWCSUM>
  ether 00:0c:29:f5:96:9f
  inet 192.168.71.3 netmask 0xffffffff broadcast 192.168.71.255
  inet6 fe80::20c:29ff:fef5:969f%em1 prefixlen 64 scopeid 0x2
  inet 192.168.71.1 netmask 0xffffffff broadcast 192.168.71.255 vhid 72
  inet 192.168.71.5 netmask 0xffffffff broadcast 192.168.71.255 vhid 72
  nd6 options=21<PERFORMNUD,AUTO_LINKLOCAL>
  media: Ethernet autoselect (1000baseT <full-duplex>)
  status: active
  carp: MASTER vhid 72 advbase 1 advskew 100
```

#6 - 03/04/2015 03:00 PM - Chris Buechler

- Assignee set to Ermal Luçi

#7 - 03/09/2015 01:00 PM - Ermal Luçi

- Status changed from Confirmed to Feedback

- % Done changed from 0 to 100

Applied in changeset [457d9c3275ff2b7ad691a38bdcb72e7177ff159a](#).

#8 - 03/09/2015 01:00 PM - Ermal Luçi

Applied in changeset [8896fe1cebdc97dcbeb59249f3bb2abd1601b979](#).

#9 - 03/10/2015 07:10 PM - Chris Buechler

- Status changed from Feedback to Confirmed

that didn't change the behavior. That whole code path there seems to have issues, like what I fixed in my last commit on this ticket. For IP aliases with a CARP parent, the oldvips array ends up containing something like:

```
array(1) {  
    ["192.0.2.205"]=>  
        string(23) "wan_vip200192.0.2.20532"  
}
```

Then under "Cleanup remaining old carps", as soon as it starts that foreach, \$oldvipar[interface] is only a single letter. 'w' in the above case.

#10 - 03/10/2015 07:39 PM - Phillip Davis

This looks related to commit <https://github.com/pfsense/pfsense/commit/89f171b052f8e72aed654d2a1c3d5a24e9bf9902>

That was applied only to master back on 16 Jan 2015. Among other things it moved

+function get_possible_listen_ips() from system.inc to interfaces.inc and changed the format of the array returned by that function.

I noticed this when testing with some stuff from master - see pull request <https://github.com/pfsense/pfsense/pull/1549> which fixes up a use of get_possible_listen_ips() by services_snmp.php

So it is a bit tricky to make changes in this area - using get_possible_listen_ips() on master is different to using it on 2.2 branch.

Also, the code in master, interfaces.inc, function get_possible_traffic_source_addresses() looks wrong. It calls get_possible_listen_ips(), which returns an array in the "new" format. get_possible_traffic_source_addresses() calls the returned array \$sourceips - then it proceeds to add VPN server/client entries to the \$sourceips array, but it in the "old" format.

I have no idea why the commit was only applied to master - but it will cause confusion for getting related code into 2.2 branch and it has bugs that I can see.

#11 - 03/11/2015 12:59 PM - Renato Botelho

- Assignee changed from Ermal Luçi to Renato Botelho

#12 - 03/11/2015 02:10 PM - Renato Botelho

- Status changed from Confirmed to Feedback

Applied in changeset [c14781e3c07dd9f82c0f0445eb5eed6c8fdb98ac](https://github.com/pfsense/pfsense/commit/c14781e3c07dd9f82c0f0445eb5eed6c8fdb98ac).

#13 - 03/11/2015 02:10 PM - Renato Botelho

Applied in changeset [c8a4eb4056a0a7927716830b11f22447e15a4f8f](https://github.com/pfsense/pfsense/commit/c8a4eb4056a0a7927716830b11f22447e15a4f8f).

#14 - 03/11/2015 09:03 PM - Chris Buechler

associated commits that got a typoed ticket #:
[934c88ee9535919b8b75b6e939b2a6becb9561bd](#)
[214c81026b6b13dc750ac971afce975117b6c493](#)

#15 - 03/11/2015 09:44 PM - Chris Buechler

- *Status changed from Feedback to Resolved*

fixed