pfSense - Bug #1605

DHCP Server should group known clients by interface

06/17/2011 07:11 AM - Willy Tenner

Status:	Pull Request Review	**Start date:**	06/17/2011
Priority:	Normal	**Due date:**	
Assignee:	Renato Botelho	**% Done:**	0%
Category:	DHCP Server	**Estimated time:**	0.00 hour
Target version:	2.5.0		
Affected Version:	All	**Affected Architecture:**	

Description

This is an old issue initially reported by LJ Rand in 2006 on another forum. No one has answered since those days. But the problem is still there:

My setup:
I've set up several VLANs behind the pfsense firewall (v2.0 RC2) and enabled DHCP on these. For all VLANs, I have enabled the DHCP setting "Deny unknown clients".

For three of those VLANs, I have set aside a dynamic range of IP addresses and enumerated the MACs of permitted clients. My VLANs are port-specific, rather than MAC-specific.

My issue:
If I had a laptop that was listed in pfsense DHCP server under VLAN A, but was plugged into a port assigned to VLAN B, I would have hoped that the firewall would consider that laptop is an unknown client on VLAN B and refused it DHCP service. Instead, it seems that pfsense does not care that the laptop was listed under VLAN A, and happily gives it an address from the dynamic range of VLAN B. That's tantamount to VLAN hopping, me thinks.

It seems that it is regardless in which MAC address list a DHCP client is listed.

Kind regards,
routerfreak

History

#1 - 06/17/2011 09:04 AM - Jim Pingle

- Subject changed from DHCP feature or VLAN-hopping anomaly? to DHCP Server should group known clients by interface

- Target version changed from 2.0 to Future

- Affected Version changed from 2.0 to All

That is not "vlan hopping" by any stretch of the imagination. If they are attached to a port on that VLAN and get an IP in that same VLAN that can route to things in that VLAN, they are in that VLAN, they didn't "hop" into it.

That said, there are probably better ways to handle the per-network known/unknown client declarations in recent versions of dhcpd. Newer versions support grouping hosts together, and also support allowing and denying unknown clients by using class matching.

But that kind of rewriting will take time and experimentation, and would be too large of a change to make it into 2.0 at this point.

#2 - 06/23/2011 09:22 PM - Evgeny Yurchenko

We are running dhcpd-4.2.1-P1 which supports grouping. The problem is 'host' cannot be related with 'subnet' (I tried different ways to group without success, moreover in some crazy configs dhcpd when started was explicitly saying that 'host' is global: "WARNING: Host declarations are global. They are not limited to the scope you declared them in."
There is a hack that can be relatively easy to implement.
We can specify 'wrong' IPs for interfaces where we do not want this MAC to appear doing something like this:

fixed-address 192.168.56.99,192.168.57.0;
If this host connects to interface 192.168.56.0/24 then it works normally but if user moves this host to 192.168.57.0/24 segment then the host gets 192.168.57.0 and can't do anything.
Let me know if this solution is acceptable and I can implement this.

#3 - 06/24/2011 04:03 AM - Willy Tenner

@Evgeny:
Uuh, nice hack. Accepted! But I don't think it's really easy to implement. Every time, someone creates/modifies/deletes a subnet and enables/disables the corresponding dhcp server you have to check/modify all host entries in dhcpd.conf reflecting the new subnet/dhcp situation. And yes, reading the man pages for dhcpd.conf, host entries do have always a global scope regardless where they are defined.
See also http://www.daemon-systems.org/man/dhcpd.conf.5.html
Thanks for the answer, will hear from you.

#4 - 06/24/2011 07:49 AM - Jim Pingle

That hack might work for the first client to connect in a subnet, but if two of them crossed at once it would not assign the same IP to them both, it would give the second client a valid IP since the first one would be in use in the leases database. At least in theory anyhow.

#5 - 06/24/2011 08:56 AM - Evgeny Yurchenko

Just tested with
host s_lan_0 {
hardware ethernet 08:00:27:e5:68:94;
fixed-address 192.168.56.99,192.168.57.0;
}
host s_lan_1 {
hardware ethernet 08:00:27:a6:50:bf;
fixed-address 192.168.56.98,192.168.57.0;

both hosts got 192.168.57.0 on the192.168.57.0/24 interface.
IPs from fixed-address do not appear in leases database, I think this is why the same IP can be given many times.

#6 - 06/24/2011 09:10 AM - Jim Pingle

It still does its ping test to ensure an IP is open before assignment, though I suppose on .0 in a /24 that would fail and it would assign it anyhow.

#7 - 06/24/2011 09:25 AM - Evgeny Yurchenko

This is what it does:
01:18:15.819317 IP 0.0.0.0.68 > 255.255.255.255.67: BOOTP/DHCP, Request from 08:00:27:a6:50:bf, length 300
01:18:15.819777 IP 192.168.57.254.67 > 192.168.57.0.68: BOOTP/DHCP, Reply, length 300
01:18:15.825990 ARP, Request who-has 192.168.57.0 tell 192.168.57.0, length 46

#8 - 06/24/2011 09:39 AM - Jim Pingle

Yeah that may be fine then, though the code to implement this should not use .0, but mathematically calculate the null route for the subnet on the interface. I believe the function gen_subnet() or similar does just that.

Either way I wouldn't put changes like this into 2.0, but if we can keep a patch handy or in a git fork it could be applied once 2.0 is branched.

#9 - 03/28/2016 01:38 PM - Jim Thompson

- Assignee set to Steve Beaver

assigned for eval (this thing is 5 year old)

#10 - 05/16/2019 08:42 PM - Daniel Koh

Segregation by class (assumed to be directly linked to interface) is now possible.

https://github.com/pfsense/pfsense/pull/4066.

Issue is similar to #4584 (https://redmine.pfsense.org/issues/4584).

#11 - 08/19/2019 01:14 PM - Jim Pingle

- Target version changed from Future to 2.5.0

Re-targeting due to pending PR

#12 - 08/27/2019 03:09 PM - Jim Pingle

- Status changed from New to Pull Request Review

#13 - 09/11/2019 02:45 PM - Renato Botelho

- Assignee changed from Steve Beaver to Renato Botelho

I'll work on it