

pfSense - Bug #6025

Load balancing fails when one gateway has a weight of 1 and another gateway has a weight >1

03/24/2016 02:55 PM - Jim Pingle

Status:	Assigned	Start date:	03/24/2016
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:	Multi-WAN	Estimated time:	0.00 hour
Target version:	2.5.0	Affected Architecture:	
Affected Version:	2.3		

Description

Strange one here:

Testing Load Balancing it seems there is an issue with weights. Given that there are two WANs each with one gateway, if one gateway has a weight of 1, and another gateway has a weight higher than 1 (2, 3, etc), only one of the gateways is used.

Example cases:

WAN=1, WAN2=1: OK

```
pass in quick on vmx1 route-to { (vmx3 203.0.113.1), (vmx0 198.51.100.1) } round-robin inet from 10.3.0.0/24 to any flags S/SA keep state label "USER_RULE: Default allow LAN to any rule"
```

WAN1=1, WAN2=2: Only uses one gateway (203.0.113.1):

```
pass in quick on vmx1 route-to { (vmx3 203.0.113.1), (vmx3 203.0.113.1), (vmx0 198.51.100.1) } round-robin inet from 10.3.0.0/24 to any flags S/SA keep state label "USER_RULE: Default allow LAN to any rule"
```

WAN1=2, WAN2=1: Only uses one gateway (203.0.113.1):

```
pass in quick on vmx1 route-to { (vmx3 203.0.113.1), (vmx0 198.51.100.1), (vmx0 198.51.100.1) } round-robin inet from 10.3.0.0/24 to any flags S/SA keep state label "USER_RULE: Default allow LAN to any rule"
```

If both weights are > 1, then it is OK:

WAN1=2, WAN2=2: OK

```
pass in quick on vmx1 route-to { (vmx3 203.0.113.1), (vmx3 203.0.113.1), (vmx0 198.51.100.1), (vmx0 198.51.100.1) } round-robin inet from 10.3.0.0/24 to any flags S/SA keep state label "USER_RULE: Default allow LAN to any rule"
```

WAN1=3, WAN2=2: OK

```
pass in quick on vmx1 route-to { (vmx3 203.0.113.1), (vmx3 203.0.113.1), (vmx0 198.51.100.1), (vmx0 198.51.100.1), (vmx0 198.51.100.1) } round-robin inet from 10.3.0.0/24 to any flags S/SA keep state label "USER_RULE: Default allow LAN to any rule"
```

Also of note, when the weights differ, even though the gateways have a specific order with repetition in the rule, pf seems to still flip back and forth, though the general ratio of the weights is respected. For example with WAN1=3, WAN2=2:

```
$ for i in {1..10}; do curl http://iptest.example.com/ip.php; done
Your IP Address is: 198.51.100.3
Your IP Address is: 203.0.113.3
Your IP Address is: 198.51.100.3
Your IP Address is: 198.51.100.3
Your IP Address is: 203.0.113.3
Your IP Address is: 198.51.100.3
Your IP Address is: 203.0.113.3
Your IP Address is: 198.51.100.3
Your IP Address is: 198.51.100.3
Your IP Address is: 203.0.113.3
```

If it turns out to be a deeper issue in pf, the easy fix would be to automatically multiply all weights by 2x if one is 1 and the other is >1.

History

#1 - 03/24/2016 02:56 PM - Jim Pingle

- Description updated

#2 - 03/24/2016 09:39 PM - Phillip Davis

All the rules that you quote look OK, so are you saying that the code that generates the rules is OK, but somehow the implementation at run-time in pf is not happening effectively?

#3 - 03/24/2016 09:59 PM - Jim Pingle

Phillip Davis wrote:

All the rules that you quote look OK, so are you saying that the code that generates the rules is OK, but somehow the implementation at run-time in pf is not happening effectively?

Correct, the rules look fine, but in practice it fails to work as expected at run time.

#4 - 03/25/2016 03:33 PM - Jim Pingle

- Target version changed from 2.3 to 2.3.1

Pushing this out a bit -- not a huge concern for now. If someone hits it they can easily work around it by adjusting weights in ways that were not possible before. Even the worst imbalance case on 2.2.x could only be 1:5, if someone wants that now they can set 2:10 and get the same effect.

#5 - 03/25/2016 03:33 PM - Jim Pingle

- Status changed from New to Assigned

- Assignee set to Marc Dye

#6 - 04/15/2016 10:08 PM - Chris Buechler

- Target version changed from 2.3.1 to 2.3.2

#7 - 07/06/2016 04:06 PM - Chris Buechler

- Target version changed from 2.3.2 to 2.4.0

#8 - 01/24/2017 06:08 AM - Renato Botelho

- Assignee deleted (Marc Dye)

#9 - 08/09/2017 11:51 PM - Luiz Souza

- Target version changed from 2.4.0 to 2.4.1

#10 - 10/12/2017 10:06 AM - Jim Pingle

- Target version changed from 2.4.1 to 2.4.2

#11 - 10/23/2017 12:18 PM - Jim Pingle

- Target version changed from 2.4.2 to 2.4.3

#12 - 01/16/2018 02:06 PM - Jim Pingle

- Target version changed from 2.4.3 to 2.4.4

#13 - 08/24/2018 03:16 PM - Jim Pingle

- Target version changed from 2.4.4 to 48

#14 - 03/12/2019 10:54 AM - Jim Pingle

- Target version changed from 48 to 2.5.0

#15 - 08/21/2019 10:22 AM - Jim Pingle

- Category changed from Rules / NAT to Multi-WAN