

## pfSense - Bug #9450

### Multiwan gateway group fail-over not working as expected (possible race condition)

04/03/2019 06:01 AM - nasir ahmed

<b>Status:</b>	New	<b>Start date:</b>	04/03/2019
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>		<b>% Done:</b>	0%
<b>Category:</b>	Multi-WAN	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	2.5.0	<b>Affected Architecture:</b>	All
<b>Affected Version:</b>	2.4.x		

#### Description

Multiwan gateway group fail-over not working as expected. After a link state change is triggered by dpinger (rc.gateway\_alarm is called) due to a higher priority link recovery, the rc.filter\_configure\_sync script fails to add the recovered gateway back to the gateway group because of a race condition.

The race condition:

rc.filter\_configure\_sync calls **get\_gwgroup\_members\_inner()** function in gwlb.inc, which manages gateway inclusion/exclusion. But this function relies on **get\_dpinger\_status()** which reads live gateway status from the dpinger instance monitoring the triggering gateway, and By the time **get\_dpinger\_status()** reads the current values, they may have fluctuated back to the gateway down range.

This, prevents **get\_gwgroup\_members\_inner()** from reactivation the gateway. In contrast, dpinger waits the "Alert interval" period which defaults to 1000 millisecond to check again for alarm conditions. By that time the loss and delay average values may move again out of the alarm range and dpinger may not trigger another gateway down alarm.

This bug results in pfsense and dpinger maintaining unmatched internal states for that particular gateway. In the described scenario the gateway status will be UP in dpinger and DOWN in pfsense. These unmatched states will be maintained until a new gateway event is triggered or a filter reload is called for any other reason.

As a solution I would suggest that the dpinger trigger values, should be written to a file by rc.gateway\_alarm and kept for at least "Alert interval" long, and then deleted. Further more, **get\_gwgroup\_members\_inner()** should read gateway loss and delay values from this file as long as it exists.

#### History

##### #1 - 04/03/2019 07:22 AM - Jim Pingle

- Category set to Multi-WAN
- Priority changed from High to Normal
- Target version set to 2.5.0
- Affected Version set to 2.4.x
- Affected Architecture set to All