pfSense - Bug #5408

broken TCP checksums with IPv6 and route-to/reply-to on gif interfaces

11/10/2015 05:27 PM - Chris Buechler

Status:	Resolved	**Start date:**	11/10/2015
Priority:	Very High	**Due date:**	
Assignee:	Chris Buechler	**% Done:**	0%
Category:	Operating System	**Estimated time:**	0.00 hour
Target version:	2.3		
Affected Version:	2.3	**Affected Architecture:**	

Description

TCP checksums on IPv6 traffic matching rules specifying route-to or reply-to end up with broken TCP checksums. Every packet from a given system has the same wrong TCP checksum. Change the IPv6 source IP and the checksum it uses changes, and still stays the same across all packets.

The issue doesn't exist in stock FreeBSD 10-STABLE.

It is at least somewhat hardware-specific. In VMware ESX, vmxnet3 doesn't exhibit the issue, but e1000 does. All physical hardware seems to be affected (RCC-VE, APU, and more).

Example on 172.27.44.174. 'ping6 google.com' works, 'fetch -6 http://google.com' fails. Take out the route-to from the rule:

```
pass out  route-to ( ... ) inet6 from ... keep state allow-opts label "let out anything from firew
all host itself"
```

and it works.

History

#1 - 11/10/2015 06:21 PM - Chris Buechler

Some relevant recent changes:
https://svnweb.freebsd.org/base?view=revision&revision=289703
https://svnweb.freebsd.org/base?view=revision&revision=290161

Appears 290161 needs MFCed.

#2 - 11/11/2015 04:20 AM - Renato Botelho

Kristof mentioned he is going to MFC 290161 today. After that happens I'm going to merge it into our branch and build new snaps

#3 - 11/11/2015 07:04 AM - Renato Botelho

- Status changed from Confirmed to Feedback

FYI, Kristof did the MFC at r290669. I've merged it into our FreeBSD-src repo and kicked off new builds. Could you please try new snapshots as soon as it is available?

#4 - 11/11/2015 08:24 AM - Jim Pingle

- Assignee changed from Luiz Souza to Chris Buechler

Of two affected systems here both have been fixed by the merge. Leaving open for more feedback but it looks OK to me so far.

#5 - 11/11/2015 03:08 PM - Jim Pingle

- Status changed from Feedback to Confirmed

- Assignee changed from Chris Buechler to Luiz Souza

There is still a problem here. It works for traffic from the firewall itself but not for traffic flowing through that hits a route-to when it enters the firewall.

For example, TCP connection enters LAN, hits a policy routing rule with route-to, exits a V6 WAN. No state is created when it exits the V6 WAN, so the SYN+ACK is denied re-entry. Remove the policy routing from the LAN rule then repeat the test and the state is created, traffic flows as expected.

#6 - 11/11/2015 10:54 PM - Chris Buechler

- Subject changed from broken TCP checksums with IPv6 and route-to/reply-to to broken TCP checksums with IPv6 and route-to/reply-to on gif interfaces

The original issue is still applicable with gif interfaces, they have the same broken checksum on every TCP packet. It's fixed on every non-gif scenario I've tried.

The issue JimP noted above is separate, opened #5424 for that.

#7 - 11/12/2015 01:58 AM - Chris Buechler

- Status changed from Confirmed to Feedback

this looks to have fixed the remainder of this issue.
https://github.com/pfsense/FreeBSD-src/commit/2e02b14e19fd0fe27055d4a6e11a65e76882bf5f

Renato/Luiz, FYI: I just pulled in some patches on that commit and
https://github.com/pfsense/FreeBSD-src/commit/fcb1a35e91beb27cdb14eeeff3aab781c0a9671c that jimt pointed out might be related so we could get snapshots to test with. Probably going to want to revert those when syncing up with FreeBSD.

#8 - 11/12/2015 11:21 AM - Jim Thompson

- Assignee changed from Luiz Souza to Chris Buechler

fixed here.

reassigning to cmb

#9 - 11/18/2015 05:18 PM - Chris Buechler

- Status changed from Feedback to Resolved

fixed