

pfSense - Bug #1047

Disable TSO, hardware checksum don't work for unassigned but active interfaces

11/27/2010 03:32 PM - Chris Buechler

Status:	Resolved	Start date:	11/27/2010
Priority:	Normal	Due date:	
Assignee:	Ermal Luçi	% Done:	100%
Category:	Interfaces	Estimated time:	0.00 hour
Target version:	2.2	Affected Architecture:	
Affected Version:	All		

Description

Interfaces that are unassigned but active (ex: part of a lagg, possibly VLAN parent-only) don't have TSO, hardware checksum offloading, and possibly LRO disabled.

Associated revisions

Revision 0ac206f9 - 11/30/2010 11:21 AM - Ermal Luçi

Ticket #1047. Make the flags the same for each member interface before adding them to lagg(4)/bridge(4)

Revision a9e44127 - 08/08/2013 12:54 PM - Renato Botelho

Fix #1047

- When advanced options (LRO, TSO and CSUM) changes, enable capabilities again on interfaces
- For lagg and bridge, check caps instead of encaps and enable flags when it's necessary for all members
- Take in consideration Disable (LRO, TSO, CSUM) options from system->advanced when configuring lagg and bridge

Revision a2dc7392 - 08/08/2013 01:05 PM - Renato Botelho

Fix #1047

- When advanced options (LRO, TSO and CSUM) changes, enable capabilities again on interfaces
- For lagg and bridge, check caps instead of encaps and enable flags when it's necessary for all members
- Take in consideration Disable (LRO, TSO, CSUM) options from system->advanced when configuring lagg and bridge

Revision 2535f6dc - 11/20/2014 08:33 AM - Ermal Luçi

Fixes #1047, Actually the code is trying to set flags on the parent. so allow it even for vlans since they will follow the parent. At least so seems on FreeBSD 10.

Revision 43517fcc - 11/24/2014 05:19 AM - Ermal Luçi

Fixes #1047, overhaul handling of flags for hardware offloading and make it work correctly for system_advanced page settings. Lagg is still a special case that needs a reboot.

History

#1 - 11/30/2010 10:25 AM - Ermal Luçi

- Status changed from New to Feedback

Patch committed.

#2 - 08/13/2011 11:13 PM - Chris Buechler

- Status changed from Feedback to Resolved

#3 - 07/14/2013 02:18 PM - Zeev Zalessky

- File config-FW3-20130714094321.xml added

- File ifconfig.jpg added

the bug is not fixed. i have same problem with Intel X540-T2 and X520-DA cards. attached my configuration and output of ifconfig. currently i use workaround using shellcmd package. i just run ifconfig ix(0-3) -lro using shellcmd, but i don't think that it is correct solution for production environment.

#4 - 07/14/2013 02:22 PM - Zeev Zalessky

i have the problem on pfsense 2.1 RC0 latest snapshot

#5 - 07/15/2013 03:40 AM - Zeev Zalessky

please reopen this bug

#6 - 08/08/2013 10:09 AM - Renato Botelho

- Status changed from Resolved to New

- Target version changed from 2.0 to 2.1

- Affected Version changed from 2.0 to 2.1

It still have problems, I'm working on it.

#7 - 08/08/2013 01:30 PM - Renato Botelho

- Status changed from New to Feedback

- % Done changed from 0 to 100

Applied in changeset [a9e4412797c5f8caeb6c7579279b2ecc7d5bfe96](#).

#8 - 08/08/2013 01:30 PM - Renato Botelho

Applied in changeset [a2dc7392fc3e3a2cbb6584dfcab8e02d26598000](#).

#9 - 08/11/2013 01:34 AM - Zeev Zalessky

- File advanced_networking.jpg added

- File config-pfsense.localdomain-20130811093828.xml added

Hi, LRO is working TSO4,TSO6 not working on LAGG interfaces.

ifconfig output:

```
ix0: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> metric 0 mtu 1500
```

```

options=403bb<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, JUMBO_MTU, VLAN_HWCSUM, TSO4, TSO6, VLAN_HWTSO>
ether 68:05:ca:0c:87:1e
inet6 fe80::6a05:caff:fe0c:871e%ix0 prefixlen 64 scopeid 0x1
nd6 options=3<PERFORMNUD, ACCEPT_RTADV>
media: Ethernet autoselect
status: no carrier
ix1: flags=8843<UP, BROADCAST, RUNNING, SIMPLEX, MULTICAST> metric 0 mtu 1500
options=403bb<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, JUMBO_MTU, VLAN_HWCSUM, TSO4, TSO6, VLAN_HWTSO>
ether 68:05:ca:0c:87:1e
inet6 fe80::6a05:caff:fe0c:871f%ix1 prefixlen 64 scopeid 0x2
nd6 options=3<PERFORMNUD, ACCEPT_RTADV>
media: Ethernet autoselect
status: no carrier
bce0: flags=8802<BROADCAST, SIMPLEX, MULTICAST> metric 0 mtu 1500
options=c01bb<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, JUMBO_MTU, VLAN_HWCSUM, TSO4, VLAN_HWTSO, LINKSTATE>
ether d8:d3:85:b9:f4:14
media: Ethernet autoselect
bcel: flags=8843<UP, BROADCAST, RUNNING, SIMPLEX, MULTICAST> metric 0 mtu 1500
options=c00bb<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, JUMBO_MTU, VLAN_HWCSUM, VLAN_HWTSO, LINKSTATE>
ether d8:d3:85:b9:f4:16
inet 192.168.100.1 netmask 0xfffff00 broadcast 192.168.100.255
inet6 fe80::dad3:85ff:feb9:f416%bcel prefixlen 64 scopeid 0x4
nd6 options=1<PERFORMNUD>
media: Ethernet autoselect (1000baseT <full-duplex>)
status: active
ix2: flags=8843<UP, BROADCAST, RUNNING, SIMPLEX, MULTICAST> metric 0 mtu 1500
options=403bb<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, JUMBO_MTU, VLAN_HWCSUM, TSO4, TSO6, VLAN_HWTSO>
ether a0:36:9f:1d:d5:00
inet6 fe80::a236:9fff:feld:d500%ix2 prefixlen 64 scopeid 0x5
nd6 options=3<PERFORMNUD, ACCEPT_RTADV>
media: Ethernet autoselect (10Gbase-Twinax <full-duplex>)
status: active
ix3: flags=8843<UP, BROADCAST, RUNNING, SIMPLEX, MULTICAST> metric 0 mtu 1500
options=403bb<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, JUMBO_MTU, VLAN_HWCSUM, TSO4, TSO6, VLAN_HWTSO>
ether a0:36:9f:1d:d5:00
inet6 fe80::a236:9fff:feld:d502%ix3 prefixlen 64 scopeid 0x6
nd6 options=3<PERFORMNUD, ACCEPT_RTADV>
media: Ethernet autoselect (10Gbase-Twinax <full-duplex>)
status: active
enc0: flags=0<> metric 0 mtu 1536
pflog0: flags=100<PROMISC> metric 0 mtu 33144
lo0: flags=8049<UP, LOOPBACK, RUNNING, MULTICAST> metric 0 mtu 16384
options=3<RXCSUM, TXCSUM>
inet 127.0.0.1 netmask 0xff000000
inet6 ::1 prefixlen 128
inet6 fe80::1%lo0 prefixlen 64 scopeid 0x9
nd6 options=3<PERFORMNUD, ACCEPT_RTADV>
pfsync0: flags=0<> metric 0 mtu 1460
syncpeer: 224.0.0.240 maxupd: 128 syncok: 1
lagg0: flags=8843<UP, BROADCAST, RUNNING, SIMPLEX, MULTICAST> metric 0 mtu 1500
options=403bb<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, JUMBO_MTU, VLAN_HWCSUM, TSO4, TSO6, VLAN_HWTSO>
ether 68:05:ca:0c:87:1e
inet6 fe80::6a05:caff:fe0c:871e%lagg0 prefixlen 64 scopeid 0xb
nd6 options=3<PERFORMNUD, ACCEPT_RTADV>
media: Ethernet autoselect
status: no carrier
laggproto lacp
laggport: ix1 flags=0<>
laggport: ix0 flags=0<>
lagg1: flags=8843<UP, BROADCAST, RUNNING, SIMPLEX, MULTICAST> metric 0 mtu 1500
options=403bb<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, JUMBO_MTU, VLAN_HWCSUM, TSO4, TSO6, VLAN_HWTSO>
ether a0:36:9f:1d:d5:00
inet6 fe80::a236:9fff:feld:d500%lagg1 prefixlen 64 scopeid 0xc
inet 10.165.128.211 netmask 0xfffff00 broadcast 10.165.128.255
nd6 options=1<PERFORMNUD>
media: Ethernet autoselect
status: active
laggproto lacp
laggport: ix3 flags=1c<ACTIVE, COLLECTING, DISTRIBUTING>
laggport: ix2 flags=1c<ACTIVE, COLLECTING, DISTRIBUTING>

```

#10 - 09/03/2013 12:34 AM - Chris Buechler

- Target version changed from 2.1 to 2.2
- Affected Version changed from 2.1 to All

#11 - 09/01/2014 11:57 AM - Jim Thompson

- Assignee set to Chris Buechler

#12 - 10/10/2014 04:00 PM - Chris Buechler

- Status changed from Feedback to New
- Assignee deleted (Chris Buechler)

This is an issue on 2.2. Example at 172.27.32.125, see its igb1 where igb1_vlan10 is assigned but igb1 isn't.

#13 - 11/19/2014 02:35 AM - Ermal Luçi

- Affected Documentation 0 added

VLANs are skipped by flags settings.

Normally a reboot should apply the right thing to the interfaces.

Reconfiguration is to be seen, in terms of applying these settings or we should force reboot when these are changed.

#14 - 11/19/2014 10:52 AM - Chris Buechler

- Status changed from New to Confirmed
- % Done changed from 100 to 0

In the circumstance described here, where the interface is in use but not directly assigned (so part of a LAGG, or a VLAN parent interface only), it's never applied to the affected interface. For instance, create lagg0 with em1 and em2 as members. em1 and em2 don't have TSO or hardware checksum set according to the system's config. Or create say em3_vlan10 and set the same there, then it's not applied to em3. Confirmed this again on latest snapshot.

#15 - 11/20/2014 08:41 AM - Ermal Luçi

- Status changed from Confirmed to Feedback

Can you try with newer snapshots?

For lagg even here there should be a specific special case since it has to go through each member.

Though reboot handles lagg quite OK so i do not think at this stage is safe to change this.

Just put it on wiki that if you want changed flags on lagg need to reboot.

#16 - 11/20/2014 08:50 AM - Ermal Luçi

- % Done changed from 0 to 100

Applied in changeset [2535f6dcc2d0898f1c89c7c0a2606c95b3f59320](#).

#17 - 11/21/2014 04:57 PM - Chris Buechler

- Status changed from Feedback to Confirmed
- Assignee set to Ermal Luçi

- % Done changed from 100 to 0

reboot doesn't handle it correctly either. There is a clear test case on 172.27.32.125 with its lagg0 and members igb2 and igb3. Please test there, Ermal. That's a test box you can do anything you need to with any time.

#18 - 11/24/2014 05:13 AM - Ermal Luçi

- Status changed from Confirmed to Feedback

Fixed Chris, though lagg still needs reboot to work on my vms and your test setup.

#19 - 11/24/2014 05:30 AM - Ermal Luçi

- % Done changed from 0 to 100

Applied in changeset [43517fcc1b616b7443b26247dc59dbd65bde2819](#).

#20 - 11/24/2014 05:47 PM - Chris Buechler

- Status changed from Feedback to Resolved

fixed

Files

config-FW3-20130714094321.xml	85.1 KB	07/14/2013	Zeev Zalesky
ifconfig.jpg	305 KB	07/14/2013	Zeev Zalesky
advanced_networking.jpg	189 KB	08/11/2013	Zeev Zalesky
config-pfsense.localdomain-20130811093828.xml	20 KB	08/11/2013	Zeev Zalesky