

## pfSense - Bug #5721

### ALTQ fails with "bandwidth for q... higher than interface" in some circumstance

12/31/2015 05:56 AM - Nikolai Fiedler

<b>Status:</b>	Resolved	<b>Start date:</b>	12/31/2015
<b>Priority:</b>	Very High	<b>Due date:</b>	
<b>Assignee:</b>	Luiz Souza	<b>% Done:</b>	100%
<b>Category:</b>	Traffic Shaper (ALTQ)	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	2.3	<b>Affected Architecture:</b>	amd64
<b>Affected Version:</b>	2.2.x		

#### Description

After the Upgrade from 2.2.5-RELEASE (amd64) to 2.2.6-RELEASE I get the following when I reboot my box.

```
There were error(s) loading the rules: /tmp/rules.debug:52: errors in queue definition - The line in question reads [52]: queue
qInternet on re2 bandwidth 150000Kb hfsc ( ecn , linkshare 150000Kb , upperlimit 150000Kb ) { qACK, qP2P, qVoIP }
There were error(s) loading the rules: /tmp/rules.debug:52: errors in queue definition - The line in question reads [52]: queue
qInternet on re0 bandwidth 150000Kb hfsc ( ecn , linkshare 150000Kb , upperlimit 150000Kb ) { qACK, qP2P, qVoIP }
```

This is the current configuration from 2.2.5-RELEASE (amd64)

```
_
altq on re0 hfsc queue { qLink, qInternet }
queue qLink on re0 bandwidth 20% qlimit 500 hfsc ( ecn , default )
queue qInternet on re0 bandwidth 150000Kb hfsc ( ecn , linkshare 150000Kb , upperlimit 150000Kb ) { qACK, qP2P, qVoIP }
queue qACK on re0 bandwidth 19.59% hfsc ( ecn , linkshare 19.59% )
queue qP2P on re0 bandwidth 4.8975% hfsc ( ecn , linkshare 4.8975% , upperlimit 4.8975% )
queue qVoIP on re0 bandwidth 32Kb hfsc ( ecn , realtime 3Mb )

altq on re2 hfsc queue { qLink, qInternet }
queue qLink on re2 bandwidth 20% qlimit 500 hfsc ( ecn , default )
queue qInternet on re2 bandwidth 150000Kb hfsc ( ecn , linkshare 150000Kb , upperlimit 150000K
b ) { qACK, qP2P, qVoIP }
queue qACK on re2 bandwidth 19.59% hfsc ( ecn , linkshare 19.59% )
queue qP2P on re2 bandwidth 4.8975% hfsc ( ecn , linkshare 4.8975% , upperlimit 4.8975% )
queue qVoIP on re2 bandwidth 32Kb hfsc ( ecn , realtime 3Mb )

altq on re1 hfsc bandwidth 5000Kb queue { qACK, qDefault, qP2P, qVoIP }
queue qACK on re1 bandwidth 19.874% hfsc ( ecn , linkshare 19.874% )
queue qDefault on re1 bandwidth 9.937% hfsc ( ecn , default )
queue qP2P on re1 bandwidth 4.9685% hfsc ( ecn , linkshare 4.9685% , upperlimit 4.9685% )
queue qVoIP on re1 bandwidth 32Kb hfsc ( ecn , realtime 1Mb ) _
```

#### Associated revisions

##### Revision 2638ddec - 03/22/2016 11:19 PM - Luiz Souza

Force the existence of bandwidth on root queue.

If the bandwidth is not set the speed of root queue is set by media status.

In this case a ruleset which contains queues with speeds higher than 100Mb/s will fail to be loaded when the root interface is connected to a fast ethernet port (100Mb/s) - but works fine on a gigabit ethernet port.

With the root queue bandwidth, sum the bandwidth of all subqueues to make sure we do not produce a broken ruleset.

Ticket #5721

**Revision ab824cde - 03/23/2016 08:57 AM - Jim Pingle**

Band-aid to stop this from breaking the firewall. Ticket #5721

**Revision 21b1367f - 03/23/2016 11:59 AM - Luiz Souza**

Pass the correct variable to get\_interface\_bandwidth(). Fixes the boot crash reported by JimP.

Spotted by: Jorge Oliveira  
Ticket #5721

**Revision 4bf6e70f - 03/24/2016 04:14 PM - Luiz Souza**

Disable the check of subqueues bandwidth until it is fixed.

Ticket #5721

**Revision 45eeb038 - 03/28/2016 05:13 PM - Luiz Souza**

Fix the sum of child bandwidth.

Now percentages are correctly handled.

Remove commented and unused old code supposed to do the same.

Ticket #5721

**Revision f6742d9a - 03/28/2016 05:13 PM - Luiz Souza**

Fixes the shaper wizards to split the bandwidth correctly.

With P2P catchall disabled, the qLink queue (using 20% of root bandwidth) was not taken into account.

Decrease 20% of qInternet bandwidth queue in this case.

Ticket #5721

**Revision 22752ae7 - 03/28/2016 05:13 PM - Luiz Souza**

Set the ALTQ root queue bandwidth.

Whenever possible set it to 1Gb, if the sum of child queues is higher than 1Gb, set it the correct value so we do not break any existing setup.

Ticket #5721

## History

---

**#1 - 12/31/2015 06:08 AM - Jim Thompson**

- Assignee set to Jim Pingle

**#2 - 12/31/2015 06:49 AM - Jim Pingle**

- Status changed from New to Feedback

Do the rules load OK after the reboot has completed? Try from Status > Filter Reload, and from the shell with "pfctl -f /tmp/rules.debug" If you get errors there, put them all here on the ticket.

If it fails during bootup but not after, it may have something to do with the NIC link state at the time not having a link with high enough bandwidth to satisfy the size of those queues.

**#3 - 12/31/2015 11:10 AM - Nikolai Fiedler**

After the upgrade to 2.2.6-RELEASE is the error back.

```
Dec 31 18:00:54 php-fpm (65802): /rc.filter_configure_sync: New alert found: There were error(s) loading the rules: /tmp/rules.debug:54: errors in queue definition - The line in question reads [54]: queue qInternet on re2 bandwidth 150000Kb hfsc ( ecn , linkshare 150000Kb , upperlimit 150000Kb ) { qACK, qP2P, qVoIP }
```

```
Dec 31 18:00:53 check_reload_status: Reloading filter
```

Here is the result from shell:

```
$ pfctl -f /tmp/rules.debug
bandwidth for qInternet higher than interface
/tmp/rules.debug:54: errors in queue definition
parent qInternet not found for qACK
/tmp/rules.debug:55: errors in queue definition
parent qInternet not found for qP2P
/tmp/rules.debug:56: errors in queue definition
parent qInternet not found for qVoIP
/tmp/rules.debug:57: errors in queue definition
pfctl: Syntax error in config file: pf rules not loaded
```

**#4 - 12/31/2015 11:13 AM - Jim Pingle**

Are the queue definitions in /tmp/rules.debug identical on 2.2.6?

"bandwidth for qInternet higher than interface" is likely the problem. Check the output of "ifconfig -a" and see what speed it reports for the interface. You have a 150Mbit/s bandwidth entered (150000Kbit/s) so if the NIC is linked at 100Mbit/s that would fail.

**#5 - 12/31/2015 11:32 AM - Nikolai Fiedler**

I have compared the queue definitions in /tmp/rules.debug from v2.2.5 and v2.2.6. The files are identical.

```

set limit table-entries 2000000
set optimization normal
set limit states 402000
set limit src-nodes 402000

#System aliases

loopback = "{ lo0 }"
WAN = "{ re1 }"
LAN = "{ re0 }"
LAN2 = "{ re2 }"

#SSH Lockout Table
table <sshlockout> persist
table <webConfiguratorlockout> persist
#Snort tables
table <snort2c>
table <virusprot>
table <bogons> persist file "/etc/bogons"
table <negate_networks>

# User Aliases
table <FritzBox> { 192.168.178.2 }
FritzBox = "<FritzBox>"
table <HL2250DN> { 192.168.178.3 }
HL2250DN = "<HL2250DN>"
table <pfsense_mooo_com> { 109.90.104.219/32 }
pfsense_mooo_com = "<pfsense_mooo_com>"
table <pfBlockerNGSuppress> persist
pfBlockerNGSuppress = "<pfBlockerNGSuppress>"

# Gateways
GWWAN_DHCP = " route-to ( re1 39.16.59.1) "

set loginterface re0

set skip on pfsync0

scrub on $WAN all fragment reassemble
scrub on $LAN all fragment reassemble
scrub on $LAN2 all fragment reassemble

altq on re0 hfsc queue { qLink, qInternet }
queue qLink on re0 bandwidth 20% qlimit 500 hfsc ( ecn , default )
queue qInternet on re0 bandwidth 150000Kb hfsc ( ecn , linkshare 150000Kb , upperlimit 150000Kb ) { qACK
K, qP2P, qVoIP }
queue qACK on re0 bandwidth 19.59% hfsc ( ecn , linkshare 19.59% )
queue qP2P on re0 bandwidth 4.8975% hfsc ( ecn , linkshare 4.8975% , upperlimit 4.8975% )
queue qVoIP on re0 bandwidth 32Kb hfsc ( ecn , realtime 3Mb )

altq on re2 hfsc queue { qLink, qInternet }
queue qLink on re2 bandwidth 20% qlimit 500 hfsc ( ecn , default )
queue qInternet on re2 bandwidth 150000Kb hfsc ( ecn , linkshare 150000Kb , upperlimit 150000Kb ) { qACK
K, qP2P, qVoIP }
queue qACK on re2 bandwidth 19.59% hfsc ( ecn , linkshare 19.59% )
queue qP2P on re2 bandwidth 4.8975% hfsc ( ecn , linkshare 4.8975% , upperlimit 4.8975% )
queue qVoIP on re2 bandwidth 32Kb hfsc ( ecn , realtime 3Mb )

altq on re1 hfsc bandwidth 5000Kb queue { qACK, qDefault, qP2P, qVoIP }
queue qACK on re1 bandwidth 19.874% hfsc ( ecn , linkshare 19.874% )
queue qDefault on re1 bandwidth 9.937% hfsc ( ecn , default )
queue qP2P on re1 bandwidth 4.9685% hfsc ( ecn , linkshare 4.9685% , upperlimit 4.9685% )
queue qVoIP on re1 bandwidth 32Kb hfsc ( ecn , realtime 1Mb )

no nat proto carp
no rdr proto carp
nat-anchor "natearly/*"
nat-anchor "natrules/*"

# Outbound NAT rules (manual)
nat on $WAN from 192.168.178.0/24 to any port 500 -> 39.16.59.126/32 static-port
nat on $WAN from 192.168.178.0/24 to any -> 39.16.59.126/32 port 1024:65535
nat on $WAN from 10.10.10.0/24 to any port 500 -> 39.16.59.126/32 static-port
nat on $WAN from 10.10.10.0/24 to any -> 39.16.59.126/32 port 1024:65535
nat on $WAN from 127.0.0.0/8 to any -> 39.16.59.126/32 port 1024:65535

```

```

nat on $WAN from 127.0.0.0/8 to any port 500 -> 39.16.59.126/32 static-port
nat on $WAN from 127.0.0.0/8 to any -> 39.16.59.126/32 port 1024:65535
nat on $WAN from 10.50.50.10/29 to any port 500 -> 39.16.59.126/32 static-port
nat on $WAN from 10.50.50.10/29 to any -> 39.16.59.126/32 port 1024:65535

# Outbound NAT rules (automatic)

# Subnets to NAT
tonatsubnets = "{ 127.0.0.0/8 192.168.178.0/24 10.10.10.0/24 }"
nat on $WAN from $tonatsubnets to any port 500 -> 39.16.59.126/32 static-port
nat on $WAN from $tonatsubnets to any -> 39.16.59.126/32 port 1024:65535

# Load balancing anchor
rdr-anchor "relayd/*"
# TFTP proxy
rdr-anchor "tftp-proxy/*"
# UPnPd rdr anchor
rdr-anchor "miniupnpd"

anchor "relayd/*"
anchor "openvpn/*"
anchor "ipsec/*"
# Allow IPv6 on loopback
pass in quick on $loopback inet6 all tracker 1000000001 label "pass IPv6 loopback"
pass out quick on $loopback inet6 all tracker 1000000002 label "pass IPv6 loopback"
# Block all IPv6
block in log quick inet6 all tracker 1000000003 label "Block all IPv6"
block out log quick inet6 all tracker 1000000004 label "Block all IPv6"
# block IPv4 link-local. Per RFC 3927, link local "MUST NOT" be forwarded by a routing device,
# and clients "MUST NOT" send such packets to a router. FreeBSD won't route 169.254./16, but
# route-to can override that, causing problems such as in redmine #2073
block in log quick from 169.254.0.0/16 to any tracker 1000000101 label "Block IPv4 link-local"
block in log quick from any to 169.254.0.0/16 tracker 1000000102 label "Block IPv4 link-local"
#-----
# default deny rules
#-----
block in log inet all tracker 1000000103 label "Default deny rule IPv4"
block out log inet all tracker 1000000104 label "Default deny rule IPv4"
block in log inet6 all tracker 1000000105 label "Default deny rule IPv6"
block out log inet6 all tracker 1000000106 label "Default deny rule IPv6"

# IPv6 ICMP is not auxiliary, it is required for operation
# See man icmp6(4)
# 1    unreachable    Destination unreachable
# 2    toobig         Packet too big
# 128  echoreq        Echo service request
# 129  echorep        Echo service reply
# 133  routersol      Router solicitation
# 134  routeradv      Router advertisement
# 135  neighborsol    Neighbor solicitation
# 136  neighboradv    Neighbor advertisement
pass quick inet6 proto ipv6-icmp from any to any icmp6-type {1,2,135,136} tracker 1000000107 keep state

# Allow only bare essential icmpv6 packets (NS, NA, and RA, echoreq, echorep)
pass out quick inet6 proto ipv6-icmp from fe80::/10 to fe80::/10 icmp6-type {129,133,134,135,136} tracker 1000000108 keep state
pass out quick inet6 proto ipv6-icmp from fe80::/10 to ff02::/16 icmp6-type {129,133,134,135,136} tracker 1000000109 keep state
pass in quick inet6 proto ipv6-icmp from fe80::/10 to fe80::/10 icmp6-type {128,133,134,135,136} tracker 1000000110 keep state
pass in quick inet6 proto ipv6-icmp from ff02::/16 to fe80::/10 icmp6-type {128,133,134,135,136} tracker 1000000111 keep state
pass in quick inet6 proto ipv6-icmp from fe80::/10 to ff02::/16 icmp6-type {128,133,134,135,136} tracker 1000000112 keep state

# We use the mighty pf, we cannot be fooled.
block log quick inet proto { tcp, udp } from any port = 0 to any tracker 1000000113 label "Block traffic from port 0"
block log quick inet proto { tcp, udp } from any to any port = 0 tracker 1000000114 label "Block traffic to port 0"
block log quick inet6 proto { tcp, udp } from any port = 0 to any tracker 1000000115 label "Block traffic from port 0"
block log quick inet6 proto { tcp, udp } from any to any port = 0 tracker 1000000116 label "Block traffic to port 0"

```

```

# Snort package
block log quick from <snort2c> to any tracker 1000000117 label "Block snort2c hosts"
block log quick from any to <snort2c> tracker 1000000118 label "Block snort2c hosts"

# SSH lockout
block in log quick proto tcp from <sshlockout> to (self) port 40022 tracker 1000000301 label "sshlockout"

# webConfigurator lockout
block in log quick proto tcp from <webConfiguratorlockout> to (self) port 443 tracker 1000000351 label "webConfiguratorlockout"
block in log quick from <virusprot> to any tracker 1000000400 label "virusprot overload table"
# block bogon networks (IPv4)
# http://www.cymru.com/Documents/bogon-bn-nonagg.txt
block in log quick on $WAN from <bogons> to any tracker 1000001561 label "block bogon IPv4 networks from WAN"
antispoof log for $WAN tracker 1000001570
# block anything from private networks on interfaces with the option set
block in log quick on $WAN from 10.0.0.0/8 to any tracker 1000001581 label "Block private networks from WAN block 10/8"
block in log quick on $WAN from 127.0.0.0/8 to any tracker 1000001582 label "Block private networks from WAN block 127/8"
block in log quick on $WAN from 172.16.0.0/12 to any tracker 1000001583 label "Block private networks from WAN block 172.16/12"
block in log quick on $WAN from 192.168.0.0/16 to any tracker 1000001584 label "Block private networks from WAN block 192.168/16"
block in log quick on $WAN from fc00::/7 to any tracker 1000001585 label "Block ULA networks from WAN block fc00::/7"
# allow our DHCP client out to the WAN
pass in on $WAN proto udp from any port = 67 to any port = 68 tracker 1000001591 label "allow dhcp client out WAN"
pass out on $WAN proto udp from any port = 68 to any port = 67 tracker 1000001592 label "allow dhcp client out WAN"
# Not installing DHCP server firewall rules for WAN which is configured for DHCP.
antispoof log for $LAN tracker 1000002620
# allow access to DHCP server on LAN
pass in quick on $LAN proto udp from any port = 68 to 255.255.255.255 port = 67 tracker 1000002641 label "allow access to DHCP server"
pass in quick on $LAN proto udp from any port = 68 to 192.168.178.1 port = 67 tracker 1000002642 label "allow access to DHCP server"
pass out quick on $LAN proto udp from 192.168.178.1 port = 67 to any port = 68 tracker 1000002643 label "allow access to DHCP server"
antispoof log for $LAN2 tracker 1000003670
# allow access to DHCP server on LAN2
pass in quick on $LAN2 proto udp from any port = 68 to 255.255.255.255 port = 67 tracker 1000003691 label "allow access to DHCP server"
pass in quick on $LAN2 proto udp from any port = 68 to 10.10.10.1 port = 67 tracker 1000003692 label "allow access to DHCP server"
pass out quick on $LAN2 proto udp from 10.10.10.1 port = 67 to any port = 68 tracker 1000003693 label "allow access to DHCP server"

# loopback
pass in on $loopback inet all tracker 1000003711 label "pass IPv4 loopback"
pass out on $loopback inet all tracker 1000003712 label "pass IPv4 loopback"
pass in on $loopback inet6 all tracker 1000003713 label "pass IPv6 loopback"
pass out on $loopback inet6 all tracker 1000003714 label "pass IPv6 loopback"
# let out anything from the firewall host itself and decrypted IPsec traffic
pass out inet all keep state allow-opts tracker 1000003715 label "let out anything IPv4 from firewall host itself"
pass out inet6 all keep state allow-opts tracker 1000003716 label "let out anything IPv6 from firewall host itself"
pass out route-to ( rel 39.16.59.1) from 39.16.59.126 to !37.24.76.0/22 tracker 1000003811 keep state allow-opts label "let out anything from firewall host itself"
# make sure the user cannot lock himself out of the webConfigurator or SSH
pass in quick on re0 proto tcp from any to (re0) port { 443 40022 } tracker 1000004121 keep state label "anti-lockout rule"

# User-defined rules follow

anchor "userrules/*"
match on { rel } proto udp from any to any tracker 1422476034 queue (qVoIP) label "USER_RULE: DiffServ/Lowdelay/Upload"
# array key "enc0" does not exist for "IPsec VPN Access" in array: {WAN LAN LAN2 } label "USER_RULE: IPsec VPN Access"
pass in quick on $LAN inet from 192.168.178.0/24 to any tracker 1422476036 keep state label "USER_RULE: Default allow LAN to any rule"
# at the break! label "USER_RULE: Default allow LAN IPv6 to any rule"

```

```

block in quick on $LAN2 inet from 10.10.10.0/24 to 192.168.178.0/24 tracker 1422476039 label "USER_RULE: B
lock access to LAN"
pass in quick on $LAN2 inet from any to any tracker 1422476040 keep state label "USER_RULE: Allow LAN 2"
# at the break! label "USER_RULE: Default allow LAN2 IPv6 to any rule"
pass in quick on $LAN2 inet proto tcp from 10.10.10.0/24 to any tracker 1422476042 flags S/SA keep state
label "USER_RULE: Default allow LAN2 to any rule"

# VPN Rules

anchor "tftp-proxy/*"
anchor "miniupnpd"

```

Here is the result from the shell

```

$ ifconfig -a
re0: flags=8943<UP,BROADCAST,RUNNING,PROMISC,SIMPLEX,MULTICAST> metric 0 mtu 1500
  options=8209b<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, VLAN_HWCSUM, WOL_MAGIC, LINKSTATE>
  ether 00:0d:b9:35:b1:44
  inet6 fe80::20d:b9ff:fe35:b144%re0 prefixlen 64 scopeid 0x1
  inet 192.168.178.1 netmask 0xfffff00 broadcast 192.168.178.255
  nd6 options=21<PERFORMNUD,AUTO_LINKLOCAL>
  media: Ethernet autoselect (1000baseT <full-duplex>)
  status: active
rel: flags=8943<UP,BROADCAST,RUNNING,PROMISC,SIMPLEX,MULTICAST> metric 0 mtu 1500
  options=8209b<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, VLAN_HWCSUM, WOL_MAGIC, LINKSTATE>
  ether f0:de:f1:41:01:de
  inet6 fe80::f2de:f1ff:fe41:1de%rel prefixlen 64 scopeid 0x2
  inet 37.24.79.126 netmask 0xfffffc00 broadcast 255.255.255.255
  nd6 options=21<PERFORMNUD,AUTO_LINKLOCAL>
  media: Ethernet autoselect (1000baseT <full-duplex>)
  status: active
re2: flags=8943<UP,BROADCAST,RUNNING,PROMISC,SIMPLEX,MULTICAST> metric 0 mtu 1500
  options=8209b<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, VLAN_HWCSUM, WOL_MAGIC, LINKSTATE>
  ether 00:0d:b9:35:b1:46
  inet6 fe80::20d:b9ff:fe35:b146%re2 prefixlen 64 scopeid 0x3
  inet 10.10.10.1 netmask 0xfffff00 broadcast 10.10.10.255
  nd6 options=21<PERFORMNUD,AUTO_LINKLOCAL>
  media: Ethernet autoselect (1000baseT <full-duplex>)
  status: active
pflog0: flags=100<PROMISC> metric 0 mtu 33144
pfsync0: flags=0<> metric 0 mtu 1500
  syncpeer: 224.0.0.240 maxupd: 128 defer: on
  syncok: 1
lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> metric 0 mtu 16384
  options=600003<RXCSUM, TXCSUM, RXCSUM_IPV6, TXCSUM_IPV6>
  inet 127.0.0.1 netmask 0xff000000
  inet6 ::1 prefixlen 128
  inet6 fe80::1%lo0 prefixlen 64 scopeid 0x6
  nd6 options=21<PERFORMNUD,AUTO_LINKLOCAL>
enc0: flags=0<> metric 0 mtu 1536
  nd6 options=21<PERFORMNUD,AUTO_LINKLOCAL>

```

With Best Regards from Bonn Germany  
Niko

**#6 - 12/31/2015 12:53 PM - Jim Pingle**

- Assignee changed from Jim Pingle to Renato Botelho

Link speeds appear to be OK, but somehow pf doesn't see that it's enough bandwidth. The syntax itself looks OK. Will need someone to look deeper down and see what, if anything, might have changed to affect this in pf.

Any chance you can try 2.3 to see if it works there with the same settings?

**#7 - 12/31/2015 01:22 PM - Nikolai Fiedler**

Sorry but for a try with v2.3 is no time. This is Prod environment. The changes in v2.3 are deeper and i do not have time for bug tracking. It will be great if someone take a look on this issue.

**#8 - 01/19/2016 04:27 AM - Steve Wheeler**

This does not seem to be resolved by upgrading to 2.3. See: CSG-79067

**#9 - 01/19/2016 03:17 PM - Nikolai Fiedler**

So what now? Should a go back to v2.2.5?  
Sorry but whatis CSG-79067?

**#10 - 01/24/2016 01:52 PM - Steve Wheeler**

The quickest way around this is to switch to PRIQ as the scheduler. That doesn't rely on absolute bandwidths and will load without errors. That may not work if you require fixed bandwidth limits or other features only offered by HFSC.  
CSG-79067 is an internal ticket number, in that particular case from a customer who hit this issue. I posted that so our developers can see that data and what we tried where I can't post customer data on a public bug report.

**#11 - 02/17/2016 06:04 AM - Steve Wheeler**

In every case where I have seen this I was able to get past it by setting the bandwidth manually on the root queue on each interface, usually to 1Gbps.  
It appears that some interface types do not report their link speed causing pf to calculate the other values incorrectly. This is especially true for lagg or vlan on lagg interfaces where the link speed is not reported at all.  
It's unclear why the re(4) NICs mentioned above appear to hit this since they do report link speed. Other NIC types such as igb(4) or em(4) do not seem effected.

Anyone hitting this should try setting the root queue bandwidth.

**#12 - 02/17/2016 08:29 AM - Jim Pingle**

- Status changed from Feedback to Assigned  
- Assignee changed from Renato Botelho to Luiz Souza  
- Target version set to 2.3

Digging through the pf source in our repo it appears we set the default bandwidth to 100 in cases when the speed cannot be detected. It may be time to increase that to 1000 given the current state of hardware:

[https://github.com/pfsense/FreeBSD-src/blob/9c716f4f2fda10bca9fd54293b2f4fe9de1e9198/sbin/pfctl/pfctl\\_altq.c#L270](https://github.com/pfsense/FreeBSD-src/blob/9c716f4f2fda10bca9fd54293b2f4fe9de1e9198/sbin/pfctl/pfctl_altq.c#L270)

**#13 - 02/17/2016 09:26 AM - Luiz Souza**

- Status changed from Assigned to Feedback

- % Done changed from 0 to 100

The default speed (when the NIC driver doesn't provide one) is now 1G.

**#14 - 02/17/2016 05:06 PM - Jim Pingle**

- Status changed from Feedback to Assigned

Either the patch didn't make it into the builds or it isn't having the intended effect. On a current snap built after that commit, it's still bailing on re and vlan+lagg setups when I try to use >100Mbit for the queues

**#15 - 03/12/2016 03:06 AM - Chris Buechler**

- Subject changed from After the Upgrade to 2.2.6 got "There were error(s) loading the rules: /tmp/rules.debug" to ALTQ fails with "bandwidth for q... higher than interface" in some circumstance

Anyone have a config that will reliably replicate this? I'm coming up empty on an APU, and with VLAN+lagg.

Steve Wheeler wrote:

Other NIC types such as igb(4) or em(4) do not seem effected.

The customer ticket noted above is on a SG-2440, so a system with igb. But it appears the line it's triggering on is pppoe0 in that case, so may not happen with igb NICs.

**#16 - 03/12/2016 08:12 AM - Jim Pingle**

Still happens on my APU, just run through the wizard using HFSC and plug in 900 Mbit/s for all the interfaces bandwidths. My LAN interface(s) are unplugged, which may be contributing to it, however if you enter 90Mbit/s instead of 900, it works.

You should have access to my APU if you can't replicate it with that info, ping me on chat for the details. I've left it in the broken state for the time being.

**#17 - 03/22/2016 11:40 PM - Luiz Souza**

- Status changed from Assigned to Feedback

The correct fix is specify the root queue bandwidth in GUI.

This makes the ruleset functional with any kind of media connection.

The bandwidth of root queue is now mandatory in 2.3.

**#18 - 03/23/2016 08:51 AM - Jim Pingle**

- Status changed from Feedback to Assigned
- Priority changed from Normal to Very High

That commit seems to have broken one of my test setups that I've noticed so far, in a rather fatal way (skips most of the boot-time config):

LWiP68m.png

**#19 - 03/23/2016 10:03 AM - Jim Pingle**

Also getting this when running through the wizard on a different system:

```
Catchable fatal error: Object of class hfsc_queue could not be converted to string in /etc/inc/shaper.inc on line 1763
```

Call Stack:

```

0.0002    232888    1. {main}() /usr/local/www/wizard.php:0
0.3453    1473152    2. eval('step8_stepsubmitphpaction();') /usr/local/www/wizard.php:155
0.3454    1473224    3. step8_stepsubmitphpaction() /usr/local/www/wizard.php(155) : eval()'d code:1
0.3454    1474128    4. apply_all_chosen_items() /usr/local/www/wizards/traffic_shaper_wizard_multi_all.inc:718
0.3710    1861568    5. hfsc_queue->add_queue() /usr/local/www/wizards/traffic_shaper_wizard_multi_all.inc:1072
0.3727    1866848    6. sprintf() /etc/inc/shaper.inc:1763

```

**#20 - 03/23/2016 02:10 PM - Jorge M. Oliveira**

I have made a "huge" PR that may fix the problem.  
 See: <https://github.com/pfsense/pfsense/pull/2785/files>  
 It needs extensive testing to ensure everything works.

**#21 - 03/24/2016 03:37 PM - Raul Ramos**

After making the wizard single Lan single Wan HFCS Scheduler type. Don't know if is related with this changes or not.

```

Catchable fatal error: Object of class hfsc_queue could not be converted to string in /etc/inc/shaper.inc on line 1759 Call Stack: 0.0001 235800
1. {main}() /usr/local/www/wizard.php:0 0.2224 2222672 2. eval('step8_stepsubmitphpaction();') /usr/local/www/wizard.php:155 0.2225 2222744
3. step8_stepsubmitphpaction() /usr/local/www/wizard.php(155) : eval()'d code:1 0.2225 2223656 4. apply_all_chosen_items()
/usr/local/www/wizards/traffic_shaper_wizard_dedicated.inc:654 0.2344 2625816 5. hfsc_queue->add_queue()
/usr/local/www/wizards/traffic_shaper_wizard_dedicated.inc:1004 0.2353 2631096 6. sprintf() /etc/inc/shaper.inc:1759

```

Can't create shaping manually, always have the error the sum of the child is greater than the parent (qRealTime), and is not. I tried without bandwidth, with percentages with fixed values. Loading shaping from old(2 weeks) config is a no go.

Setup:

```

WAN > qRealTime > qMax + qVoIP + qACK + qGames; WAN > qNoRealTime > qHigh + qDefault + qOtherHigh + qP2P
LAN > qRealTime > qMax + qVoIP + qACK + qGames; LAN > qNoRealTime > qHigh + qDefault + qOtherHigh + qP2P ; LAN > qLocal +
qLocalACK(this one is for pFsense traffic not take account in the internet traffic)

```

2.3-BETA (amd64) built on Thu Mar 24 05:16:15 CDT 2016 and gitsync

**#22 - 03/24/2016 04:23 PM - Luiz Souza**

This change was reverted while I'm working on a proper fix.

Sorry for the breakage.

**#23 - 03/25/2016 06:29 AM - Orsiris de Jong**

Having the same issues here, with Beta 2.3 from 22 March.  
I'd be glad to test fixes :)

**#24 - 03/28/2016 05:25 PM - Luiz Souza**

- Status changed from Assigned to Feedback

New fixes were committed to address the issues from previous commits.

Please test.

**#25 - 03/29/2016 01:38 AM - Chris Linstruth**

This was fresh pfSense-CE-memstick-ADI-2.3-BETA-amd64-20160328-1625.img.gz after uploading a 2.2.6 config.

Configuring CARP settings...done.

Syncing OpenVPN settings...done.

Configuring firewall..

Fatal error: Call to undefined method dnqueue\_class::SetRoot() in /etc/inc/shaper.inc on line 3525

Call Stack:

```
0.0075 213832 1. {main}() /etc/rc.bootup:0
90.3060 5264992 2. filter_configure_sync() /etc/rc.bootup:268
90.5585 5309480 3. filter_generate_dummynet_rules() /etc/inc/filter.inc:293
90.5585 5309616 4. read_dummynet_config() /etc/inc/shaper.inc:4706
90.5589 5317624 5. dnpipe_class->add_queue() /etc/inc/shaper.inc:4598
```

PHP ERROR: Type: 1, File: /etc/inc/shaper.inc, Line: 3525, Message: Call to undefined method dnqueue\_class::SetRoot()Starting CRON... done

Looks like something similar at 1757.

**#26 - 03/29/2016 01:54 AM - Luiz Souza**

Weird, all the calls to SetRoot() were removed from code, I'll double check. Thanks.

**#27 - 03/29/2016 03:16 AM - Orsiris de Jong**

Build 2.3-BETA (amd64) Tue Mar 29 01:50:40 CDT 2016 resolves my shapper issue with the queues sum being reported greater than the parent queue.

**#28 - 04/05/2016 12:34 AM - Chris Buechler**

*- Status changed from Feedback to Resolved*

all fixed