PF QUEUING

FOR OPENBSD 5.5 AND NEWER

HISTORY

- ALTQ (included 3.0 – 5.5)
  - Natively Off
  - Required a daemon (3.0 – 3.4)
  - 3 Schedulers
    - Class Based
    - Priority
    - HFSC

NEW QUEUING SYSTEM

- New Syntax (Easier in my opinion)
- Always On
- Priority's are set via rules
- HFSC Scheduler[1]

FUNCTIONALITY

- Only shapes **outbound** traffic on the *if*spec interface *
- Only shapes traffic when no bandwidth available *

BENEFITS

- Prevent timeouts from an overloaded outbound connection
- Maintain purchased speed on an unmetered drop (billed by hourly average)
- Prioritize Voice and Video applications and prevents jitter
- Insure small file transfers and page load times are faster than larger file transfers

MINIMUM REQUIREMENTS

- At least 1 queue specified in pf.conf to function

```
queue root on em0 bandwidth 100M default
```

USAGE SCENARIOS

- Web Server
- Home or Office Firewall
- Router

WEB SERVER

```
queue root on em0 bandwidth 100M max 100M
  queue http parent root bandwidth 30M max 30M burst 80M for 3000ms
  queue ssh parent root bandwidth 10M
    queue ssh_int parent ssh bandwidth 5M
    queue ssh_bulk parent ssh bandwidth 5M

pass in on em0 proto tcp to port {80,443} set queue http
pass in on em0 proto tcp to port ssh set queue (ssh_bulk, ssh_int)
```

HOME FIREWALL

```
queue root em0 bandwidth 10M max 10M burst 12500K for 6400ms
  queue dns parent root bandwidth 500K
  queue ack parent root bandwidth 4M
  queue bulk parent root bandwidth 5500K default
 match on em0 proto tcp to port {21,80,443,5223} set queue (bulk,ack) set prio (3, 6)
 match out on em0 proto udp set prio 4
 match out on em0 proto {tcp,udp} from any to any port 53 set queue dns set prio 7
```

ROUTER

```
queue l3 on em0 bandwidth 500M max 500M default
queue ntt on em1 bandwidth 200M max 200M default
queue att on em2 bandwidth 300M max 300M default
match on em0 set queue l3
match on em1 set queue ntt
match on em2 set quete att
```

MATH FOR BANDWIDTH CALCULATION

Physical interface speed(bps) / (frame size(B) * 8 (B to b)) = packets per second *

Frame size 84 to 1538 bytes (no vlan's or Q in Q)

```
1,000,000,000 bps / (84 * 8) = 1,488,096
1,000,000,000 bps / (1538 * 8) = 81,274
```

* For a single direction on a single interface

MATH FOR BANDWIDTH CALCULATION

TCP ACK bandwidth

Incoming PPS with largest frame size * 66 bytes * 8 = outgoing ACK space in bps

```
81,274 * 66 * 8 = 42,912,672 bps
```

UDP Space

DNS 1 to 5% of upstream

VoIP stuff into bulk or default with higher priority or carve out space by defining a queue and assign rules

Queue Space

Make sure the sum of bandwidth specified on child queues does not exceed the root queue

SPECIFYING QUEUS'S

set queue *qname* or set queue (qname , qname)

add to block, match or pass rules

Pass rules: incoming (my preference) or no direction specified

Match rules: any or no direction specified

Block rules: any direction

SETTING PRIORITIES

- PF allows priorities to be set on packets via their matching rule
- 0 – 7, higher the number higher the priority
- Default priority if none is specified is 3
- Packets with a higher priority are processed before lower priorities

PRIORITIES

- set prio # or set prio (# , #)
- You can operate with the minimum of 1 queues and rules to set priorities.

```
queue root on em0 bandwidth 10M max 10M default
match out on em0 proto tcp set prio (3,6)
match out on egress proto {tcp,udp} from any to any port 53 set prio 7
```

CAVEATS

- PF is stateful and queue's must be set when the state is created
  - does not matter which way the traffic is going
- PF will not shape traffic until there is no bandwidth left
- PF will not drop packets until the qlimit(aka qlength) is full *
- If `max` is not specified on the root queue, pf will allow the bandwidth specified to be exceeded (when bandwidth is lower than NIC speed)
- When max is set on child queues, that queue can longer borrow from unused bandwidth

CHECKING UP ON THINGS

- `systat q`
  - Will show the queue tree and it's current snapshot updated every second
- `pfctl -vvs queue`
  - Different layout more info also updated every 5 seconds

TROUBLESHOOTING

- Traffic being wrongly classified or dropped where it shouldn't be
- Add `log` to the rules that specify the queues

```
tcpdump -nettr /var/log/pflog
```

QUESTIONS

- ?

BY JIM HOFFMAN

jim (at) securebytes.org

REFERENCES

1) Hansteen, Peter N.M. *The Book of PF: A No-Nonsense Guide to the OpenBSD Firewall; 3rd Edition*. N.p.: No Starch, 2014. 123. Print.