

The two-tier Internet, delivered*

Improving NANOG Wireless with QoS

***Whether you knew it or not, you probably liked it**

What are the problems with wireless?

- **.11 mac is naturally 'half duplex'**
- **Generally pps limited [1]**
 - 479 for 11b
 - 2,336 for 11g/11a
- **Positive ACK considered inefficient**
 - .16 and friends implement NACK for ARQ
 - NACK often implies contentionless
- **Limited control of stations uplink**
 - No scheduler provided in standard .11
 - 802.11e to address this, no implementations yet
- **Only 17 to 20% of NANOG laptops appear to support .11a**
 - Alternative is b/g, rife with performance issues

[1] http://www.oreillynet.com/pub/a/wireless/2003/08/08/wireless_throughput.html

What if there's no wireless QoS?

congeztion

deliver(3)d

How was this improved?

- **Diffserv to the rescue!**
- **Map DSCP and IP PREC to various CoS levels within AP's**
- **Color specific packets at the NANOG border router**
 - **Set unmarked tcp 22 & 23 to prec 6**
 - **Set esp/ah to dscp 24 (cs3)**
- **In our case, 'Two Tier' is a lie**
- **We're using "K-th tier Internet"**

Ok, what's the config look like?

- **Define classes**

```
class-map match-any realtime
  match ip dscp ef
class-map match-any interactive
  match ip precedence 6
class-map match-any good
  match ip dscp af31
  match ip dscp cs3
class-map match-any best-effort
  match ip dscp default
```

- **Use class in policy**

```
policy-map nanog-wireless
  class realtime
    set cos 6
  class interactive
    set cos 5
  class good
    set cos 3
  class best-effort
    set cos 1
  class class-default
    set cos 0
```

Config cont.

- **Bind to interface**

```
interface Dot11Radio0
  no ip address
  service-policy output nanog-wireless      !attach policy map here
  no ip route-cache
  load-interval 30
  !
  ssid nanog-arin
  !
  speed basic-5.5 9.0 11.0 12.0 18.0 24.0 48.0 54.0      !make 5.5 mbit minimum
  no power client local
  power client 30
  power local cck 20
  power local ofdm 20
  packet retries 32 drop-packet      !fix this lil' devil
  channel 2412
  station-role root
  bridge-group 1
  bridge-group 1 subscriber-loop-control
  bridge-group 1 block-unknown-source
  no bridge-group 1 source-learning
  no bridge-group 1 unicast-flooding
  bridge-group 1 spanning-disabled
```

Router firewall filter config

```
term 20 {
  from {
    tcp-initial;
  }
  then {
    count nanog38-term20;
    forwarding-class friggin-awesome;
    accept;
  }
}
term 40 {
  from {
    protocol tcp;
    port 22;
  }
  then {
    count nanog38-term40;
    forwarding-class friggin-awesome;
    accept;
  }
}
term 50 {
  from {
    protocol tcp;
    port 23;
  }
  then {
    count nanog38-term50;
    forwarding-class friggin-awesome;
    accept;
  }
}
```

```
term 100 {
  from {
    protocol esp;
  }
  then {
    count nanog38-term100;
    forwarding-class bestest-effort;
    accept;
  }
}
term 110 {
  from {
    protocol ah;
  }
  then {
    count nanog38-term110;
    forwarding-class bestest-effort;
    accept;
  }
}
term 200 {
  then {
    count nanog38-term200;
    accept;
  }
}
```

Thanks Benson!

Router class of service config 2

```
code-point-aliases {
  inet-precedence {
    cs3 011;
    cs6 110;
  }
}
forwarding-classes {
  queue 1 friggin-awesome;
  queue 2 bestest-effort;
  queue 0 best-effort;
}
interfaces {
  fe-0/2/0 {
    unit 171 {
      rewrite-rules {
        inet-precedence nanog38-rewrite;
      }
    }
  }
}
rewrite-rules {
  inet-precedence nanog38-rewrite {
    forwarding-class friggin-awesome {
      loss-priority low code-point cs6;
      loss-priority high code-point cs6;
    }
    forwarding-class bestest-effort {
      loss-priority low code-point cs3;
      loss-priority high code-point cs3;
    }
  }
}
```


What were the results?

- **Hopefully happy users!**
 - **Please inform me to the contrary**
- **Standard UDP traceroute toward saturated .11b client**

```
12  acr1-so-0-0-0.chicago.savvis.net (208.172.3.54)  6.118 ms  6.244 ms  9.767 ms
13  scr1-at-1-2-0-960.stls1.savvis.net (208.172.1.202) 15.919 ms 12.430 ms 12.049 ms
14  savvis-internet.scr1-stls1.savvis.net (209.16.211.22) 13.333 ms 13.287 ms 18.241 ms
15  192.35.169.66 (192.35.169.66) 13.498 ms 12.744 ms 13.197 ms
16  hatebook.dhcp.nanog.merit.net (192.35.165.145) 102.148 ms 123.896 ms 151.478 ms
```

- **Tcptraceroute port 22 toward saturated .11b client**

```
12  acr1-so-0-0-0.chicago.savvis.net (208.172.3.54)  11.804 ms  6.029 ms  6.227 ms
13  scr1-at-1-2-0-960.stls1.savvis.net (208.172.1.202) 12.637 ms 66.029 ms 12.277 ms
14  savvis-internet.scr1-stls1.savvis.net (209.16.211.22) 13.007 ms 13.168 ms 12.923 ms
15  192.35.169.66 14.928 ms 19.139 ms 27.021 ms
16  hatebook.dhcp.nanog.merit.net (192.35.165.145) [closed] 14.683 ms 20.215 ms 26.944 ms
```

What's being used?

```
Class-map: realtime (match-any)
  21 packets, 5120 bytes
  30 second offered rate 0 bps, drop rate 0 bps
Match: ip dscp ef
  21 packets, 5120 bytes
  30 second rate 0 bps

Class-map: interactive (match-any)
  577039 packets, 255652271 bytes
  30 second offered rate 40000 bps, drop rate 0 bps
Match: ip precedence 6
  577039 packets, 255652271 bytes
  30 second rate 40000 bps

Class-map: good (match-any)
  15696 packets, 13442248 bytes
  30 second offered rate 110000 bps, drop rate 0 bps
Match: ip dscp af31
  0 packets, 0 bytes
  30 second rate 0 bps
Match: ip dscp cs3
  15696 packets, 13442248 bytes
  30 second rate 110000 bps

Class-map: best-effort (match-any)
  3037705 packets, 2037908959 bytes
  30 second offered rate 223000 bps, drop rate 0 bps
Match: ip dscp default
  3037705 packets, 2037908959 bytes
  30 second rate 223000 bps

Class-map: class-default (match-any)
  3648496 packets, 736524578 bytes
  30 second offered rate 28000 bps, drop rate 0 bps
Match: any
```

What's being caught?

```
admin@savvis-nanog> show firewall filter nanog38-cos
```

```
Filter: nanog38-cos
```

```
Counters:
```

Name	Bytes	Packets
nanog38-term20	38150121	76983 <- this is TCP-Initial traffic
nanog38-term110	0	0 <- this is AH traffic
nanog38-term100	54883488	5962 <- this is ESP traffic
nanog38-term50	4384	75 <- this is TCP port 23 traffic
nanog38-term40	86876026	184465 <- this is TCP port 22 traffic
nanog38-term200	1300767611	1690550 <- this is everything else

Other observations

- **12.3(8)JEA does not support WFQ**
 - All classes are FIFO queued
- **Class bandwidth also not supported**
 - Classes may starve each other
 - Must ingress police to prevent
 - Befogged by a 'variable rate' towards user stations
- **Default action is to deassociate client when max-retries reached**
 - Set to 'drop' instead (duh)
- **CoS values are hybrid 'priority queuing' knobs**
 - Also controls DCF backoff function
 - Progressively shorter 'retry' intervals for progressively higher priority classes
 - See following defaults chart

Cisco wireless CoS map

Priority	Fixed slottime Number	CWmin	CWmax
0	6	15	1023
1	2	15	1023
2	1	7	15
3	1	3	7

```
NANOG-AP-9(config-if)#traffic-class ?
 0      Parameters for priority 0
 1      Parameters for priority 1
 2      Parameters for priority 2
 3      Parameters for priority 3
 4      Parameters for priority 4
 5      Parameters for priority 5
 6      Parameters for priority 6
 7      Parameters for priority 7
background Parameters for the background access class
best-effort Parameters for the best effort access class
video   Parameters for the video access class
voice   Parameters for voice access class
```

Questions?

tk@5ninesdata.com

Thanks