Tripping on QoS

John Kristoff
jtk@ultradns.com
UltraDNS

based on work done at Northwestern University

Thanks to joe@uoregon for title inspiration
and booloo for the link to blotterart.com

Basic Picture I'm Working From 100 Mb/s Ends, Else 1 Gb/s+ On An All Cisco Network

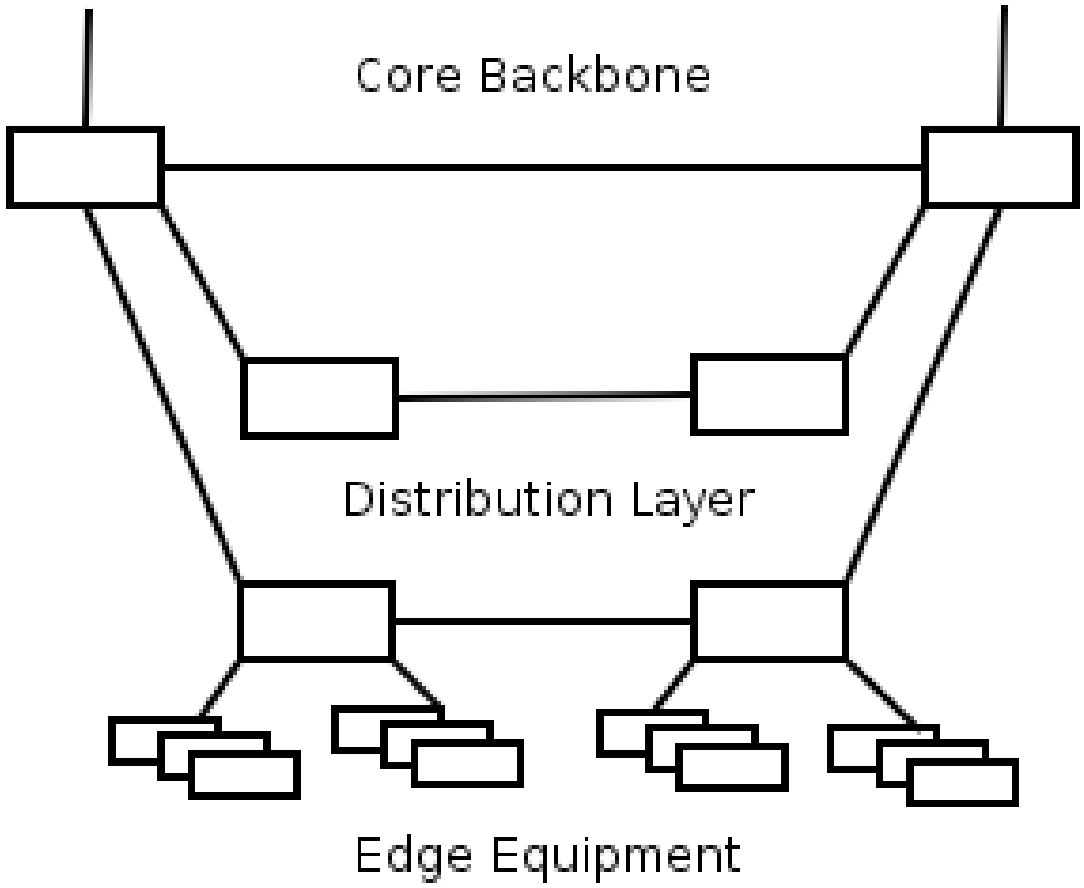

Cisco 6509 Native IOS CoPP and QoS

- Two seemingly unique feature sets are related
- You need to enable 'mls qos' for either feature
- Welcome to class-maps and policy-maps
- If you dislike ACLs, you're gonna hate these
- Show commands and counters are imperfect
- Troubleshooting w/o ACL logs or NetFlow, yay!

Control Plane Policing (CoPP)

- Single set of rules to protect the router
- Acts on packets **to** the router from any interface
- You can rate limit or filter
- Akin to a loopback filter in JUNOS
- In a nutshell, I think this is good to try to do
- ...but it can get complicated
- ...and painful if not well managed

Quality of Service (QoS)

- Or Class of Service (CoS) if you prefer
  - there is a difference, but let's be lazy
- Customizes packet queueing and forwarding
- In a nutshell, I think you usually don't need this
  - ...particularly for what 6509s are used for
  - ...but theology may trump science
- Simple or complex implementation?
  - each approach is a trade-off

CoPP Questions to Ask

- What traffic to the router do you have?
  - you are probably forgetting something
- Do you want to rate limit or just permit/deny?
- Do you want to have a default deny policy?
  - it might be easier in the long run if you did
  - it might be harder in the short run to setup
- Is there any traffic CoPP doesn't support?
  - (e.g. ARP, IPv6)

Once Bitten, Twice Shy

- Your CoPP policy might have this:
  - permit udp host 0/32 host 255.255.255.255 eq bootps
- But did you know you might also need this?
  - permit udp host dhcpd eq bootps any eq bootps
- Having a catch all rate limit can really bite you
  - ip sap listen enabled and broken CoPP config
  - http://www.uoregon.edu/~llynch/mboned/msg03010.html

Things to Consider with CoPP

- BOOTP/DHCP
- BGP/MSDP peers and PIM/OSPF neighbors
- IGMP (224/4) – need a 'no forward knob'
- SNMP/ICMP/TELNET/SSH management
- NTP/TACACS/RADIUS/DNS services
- Who sends PINGs/traceroutes (lots of people)
- Link capacity and QoS testing tools
- “The network is losing packets” emails

Know This About CoPP

- If you 'mls qos' enable and nothing else...
- All your ports are marked untrusted by default
  - which means the DiffServ field is zero'd for all packets passing through the router
  - should you care? maybe not, but it is nice to know that this is happening
- In addition, strict/RED queues are activated

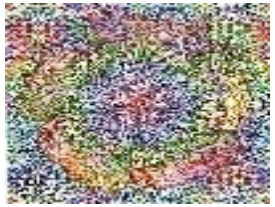

Some QoS Terms I May Use

- Differentiated Services (DiffServ)
- DiffServ Code Point (DSCP)
  - Expedited Forwarding (EF)
  - Assured Forwarding (AF) or (AFxy)
  - scavenger (CS1)
  - best effort or Default (CS0)
- Drop, mark, police, queue, shape, congestion
- Random Early Detection (RED)

A Typical QoS Approach

- Upon entering the network, traffic is classified, marked, re-marked or trusted for select endpoints
- As traffic travels into the backbone some out-of-profile traffic may be policed
- Different classes put into their respective queues towards the destination
  - voice (marked EF) always goes first
  - other flows share in some fashion

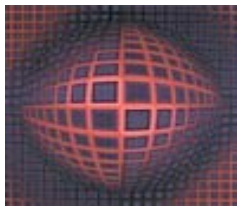

Some QoS Questions to Ask

- Trust ends to set DiffServ codepoint (DSCP)?
  - service theft vs. verification/mark complexity
- Police EF traffic? Treat all hosts/nets equally?
  - what will you break?
- How will you test this, how will you monitor it?
- Are you good at configuration management?
- Have you accounted for all edges?
- Doing this just for VoIP?

The Smart Edges Can Do a Lot

- Catalyst 3550 and greater are feature-rich
- If we see scavenger, just forward it
- If we see ACL matching VoIP traffic, mark it
- All ports edge are set to untrust by default
- All other traffic set to best effort (CS0)
- Queues pretty much at defaults or BCP

6509s Worry About Congestion

- Except direct attachments, ports set to trust
  - if we miss an edge, it is a potential problem
- 384 Kb/s EF per source address max rate
  - thought to be safe for single user VoIP ends
  - defining end user net is a bit of a gray area
- Default queues are used
- Also used some generic protocol rate limits

Know This About QoS

- Be careful about remarking portions of a flow
  - it could result in packet re-ordering
- Cisco marks HSRP as best effort (CS0)
  - it might be nice if it was some AFxy DSCP
- Careful using rate limits on untrusted ports
  - conforming traffic may retain original DSCP
  - careful configuration required

Current Status

- Not done, not fully tested/baked
- No monitoring, difficult to do in network gear
- Initial config deployment went mostly OK
  - broke some video conferencing services when per source 384 Kb/s EF rate incorrectly applied to server subnet
- Glad not to be drinking the kool-aid anymore :-)

CoPP References

- My initial report, probably don't do exactly this
  - http://aharp.ittns.northwestern.edu/papers/copp.html
  - has links to Cisco authoritative material
- Early control plane talk, using rACLs
  - http://www.nanog.org/mtg-0405/mcdowell.html

DiffServ RFCs

- RFC 2474 – Definition of Differentiated Services Field (DS Field) in the Ipv4 and Ipv6 Headers
- RFC 2745 – An Architecture for Differentiated Services
- RFC 2597 – Assured Forwarding PHB Group
- RFC 3246 – An Expedited Forwarding PHB (Per-Hop Behavior)

Cisco Documentation I

- Catalyst 3550 Switch Software Configuration Guide - Configuring QoS
- QoS Scheduling and Queueing on Catalyst 3550 Switches
- Catalyst 3750 Switch Software Configuration Guide – Configuring QoS
- Catalyst 6500 Series Software Configuraiton Guide – Configuring QoS

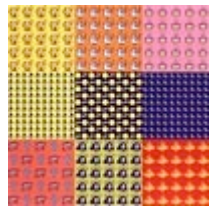

Cisco Documentation II

- DiffServ – The Scaleable End-to-end Quality of Service Model
- Quality of Service for Voice over IP
- Catalyst 6500 Series Cisco IOS Software Configuration Guide, 12.2SX – Configuring PFC QoS
- QoS Classification and Marking on Catalyst 6500/6000 Series Switches That Run Cisco IOS Software
- QoS Policing on Catalyst 6500/6000 Series Switches

Cisco Documentation III

- QoS Output Scheduling on Catalyst 6500/6000 Series Switches Running Cisco IOS System Software
- QoS Output Scheduling on Catalyst 6500/6000 Series Switches Running CatOS System Software
- Enterprise QoS Solution Reference Network Design Guide

Recent QoS Research and Deployment Reports

- ACM Workshop on Revisiting IP QoS: Why do we care, what have we learned? (RIPQOS)
- DOE NGI Testbed – Project Close Out Report, Joe Burrescia, James Leighton, 3 December 2001
- MOREnet Bandwidth Management / Quality of Service

Appendix A
Class Selector (CS) codepoints

- CS0 = 000000 (0) Default, best-effort
- CS1 = 001000 (8) Priority
- CS2 = 010000 (16) Immediate
- CS3 = 011000 (24) Flash
- CS4 = 100000 (32) Flash override
- CS5 = 101000 (40) CRITIC/ECP
- CS6 = 110000 (48) internetwork control
- CS7 = 111000 (56) network control

Appendix B
Assured Forwarding (AF) Codepoints

- AF11 = 001010 (10)
- AF12 = 001100 (12)
- AF13 = 001110 (14)
- AF21 = 010010 (18)
- AF22 = 010100 (20)
- AF23 = 010110 (22)
- AF31 = 011010 (26)
- AF32 = 011100 (28)
- AF33 = 011110 (30)
- AF41 = 100010 (34)
- AF42 = 100100 (36)
- AF43 = 100110 (38)

Appendix C
Other Codepoints

- Expedited Forwarding (EF)
  - EF = 101110 (46)
- Internet2 Qbone Scavenger Service (QBSS)
  - Scavenger = 001000 (8)
  - Note, also known as CS1