"Quality of Service", a common misconception

Julien Goodwin
jgoodwin@studio442.com.au
@laptop006

Disclaimer

- This is general information
- Tainted by personal preference
- Each slide could be a full day talk and not cover all the edge cases

What is QoS?

- Deciding which packets to drop during congestion
- No congestion == No benefit
- No ability to differentiate packets == No benefit

How does it work?

- Assign traffic to >1 “traffic class”
  - Via DSCP bits (or MPLS EXP etc.)
  - Or using ACL filters for IP src/dest/ports
- Assign each traffic class:
  - Bandwidth (Mb/sec)
  - Queue size (ms, MB, %)

What does congestion mean?

- An attempt to send data above line rate
- This is congestion
- After “some time” the packet buffers fill up and block new packets
- This is called tail dropping, and is the more critical problem
- On systems using QoS buffers are sliced per-queue reducing this time

On buffer sizes

- Some devices have buffers so small that you shouldn't enable QoS
  - Ex: Cisco 2960 series
- Some devices have buffers so large you need QoS to avoid "buffer bloat"
  - Most "carrier grade" routers (Ex: Juniper T series)
- Some devices are just horrible
  - The Cisco Catalyst 6500

General rules

- "Low latency" == small buffers, frequent drops
- "High throughput" == large buffers, few drops
- In practice you need to balance the two

Software routers

- Buffer sizes on software routers is largely a config option, look into your specific devices
- For Linux see the cerowrt project

Choosing traffic classes

- NC - Routing & switching protocol traffic
- AF “High pri”
- AF “Low pri”
- BE “Best effort” - Everything else
- “Scavenger” - Very low pri traffic, eg backups

Conclusion

QoS can be a benefit, but unless you employ good network engineers the dangers can outweigh any benefits.

Think, and test, carefully before deployment.

Questions?

Julien Goodwin
jgoodwin@studio442.com.au
@laptop006