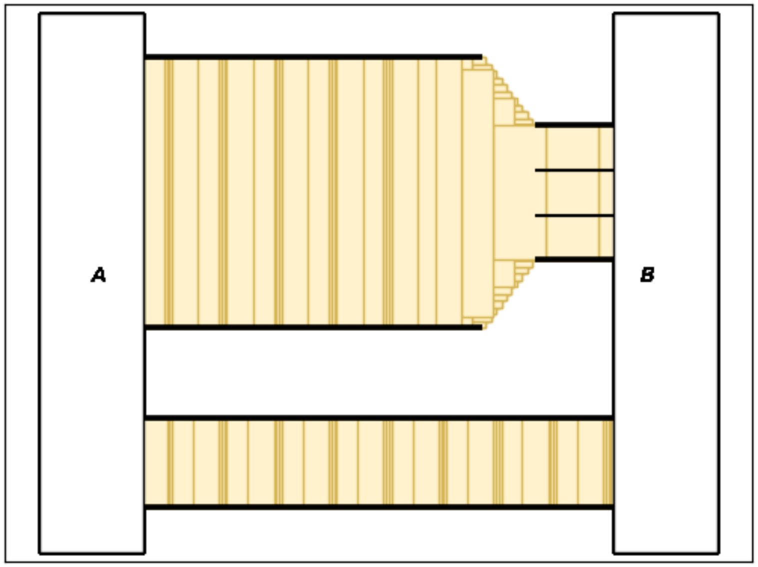

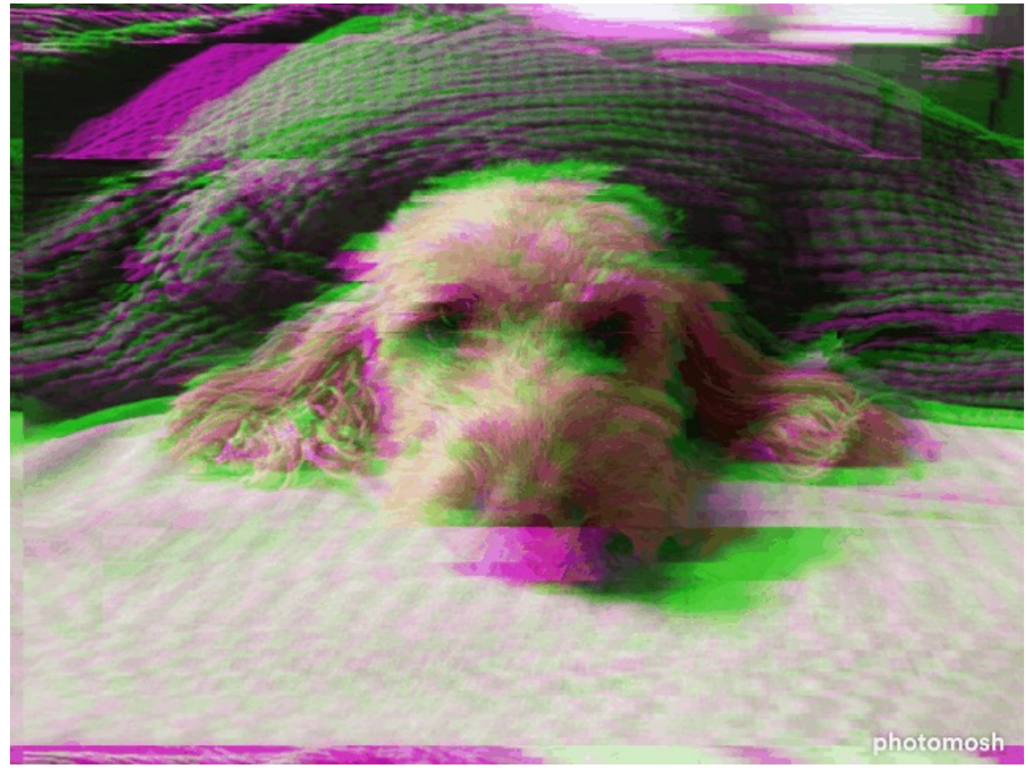

Dual Queue Networking: *How to Prepare Your Network for L4S & NQB*

NANOG 90

Presented by Jason Livingood

COMCAST

Latency: what is it & why it matters

Everything we believe about Internet performance is (mostly) wrong.

Bandwidth (marketed as speed) is important but is no longer the dominant performance factor.

Nearly all end user Internet QoE issues can be attributed to latency.

Problem is - no one knows what latency is. (or if they do – they are wrong)

From IDLE to WORKING Latency

Changing Terms, Changing Mindsets

- When you think of "latency" – think of "**delay**", "**responsiveness**" or "**response time**"
- **Idle Latency:** today's artificial measure of responsiveness when a network connection is unused
- **Working Latency:** real-world measure of responsiveness when a network connection is being used (under load)

Latency: the X-Factor of QoE

- It's a *latency* problem (not bandwidth):
  - Zoom dropping or jerky video/audio
  - Laggy/slow gaming
  - WiFi slow downs when several people are online
  - Poor VoIP quality
  - Delays loading web pages

- We used to think there was a tradeoff: EITHER high throughput OR low latency
  - <u>Turns out we can have both</u>!!

Latency: what is it & why does it matter

Not Just Gaming, AR, VR, etc.

-Latency / delay / responsiveness matters for every application ***where a user is interacting with an interface: screen, camera, microphone, haptics, etc***.

-Improving active latency improves most end user apps

-New apps will arise where delay to local storage is equivalent to delay to a network-based resource

Dual Queue Basics

Low latency networking targets “queuing delay”

(image from Nokia)

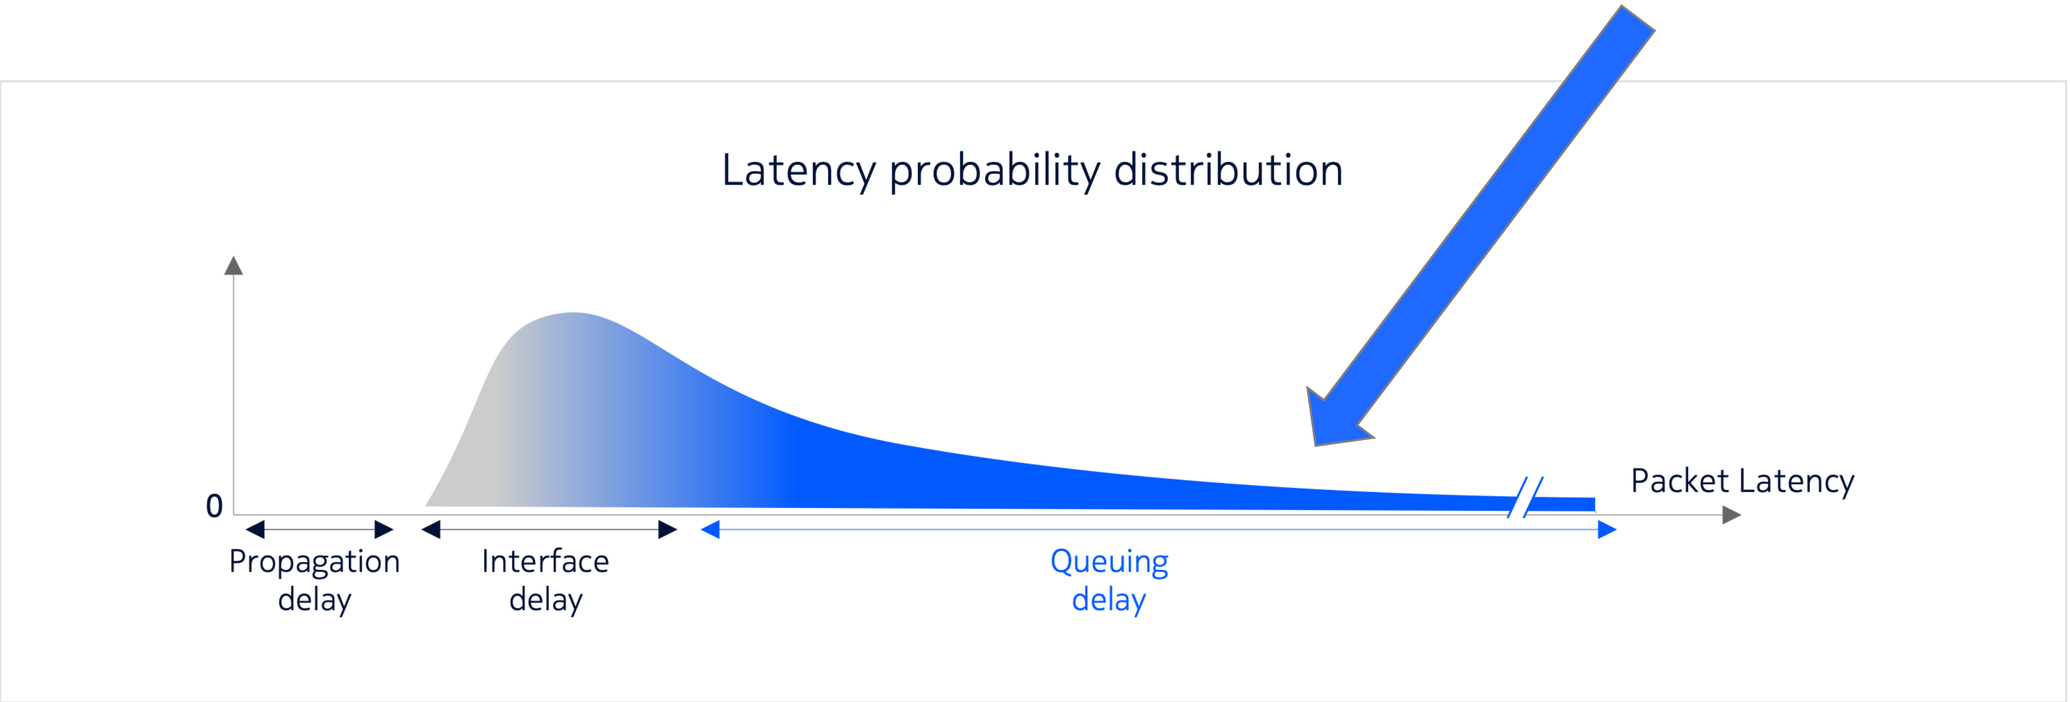

Dual Queue Basics

- There is always a bottleneck link on the end-to-end path
- Two fundamental traffic types = two separate network queues at bottlenecks
- Can be incrementally deployed – does not depend on full internet-wide deployment
- Loose coupling between layers (between network & app developer)

Deployment Principles:

- Only apps mark traffic, not the (access) network (e.g., with a DPI middlebox)
- Any app/edge provider should be able to use it
- Network operators just need to let the ECN and/or DSCP marks to flow across the network without bleaching or other modification

Dual Queue Basics

Establish 2nd network queue at bottlenecks, with a shallower queue depth

(image from Nokia)

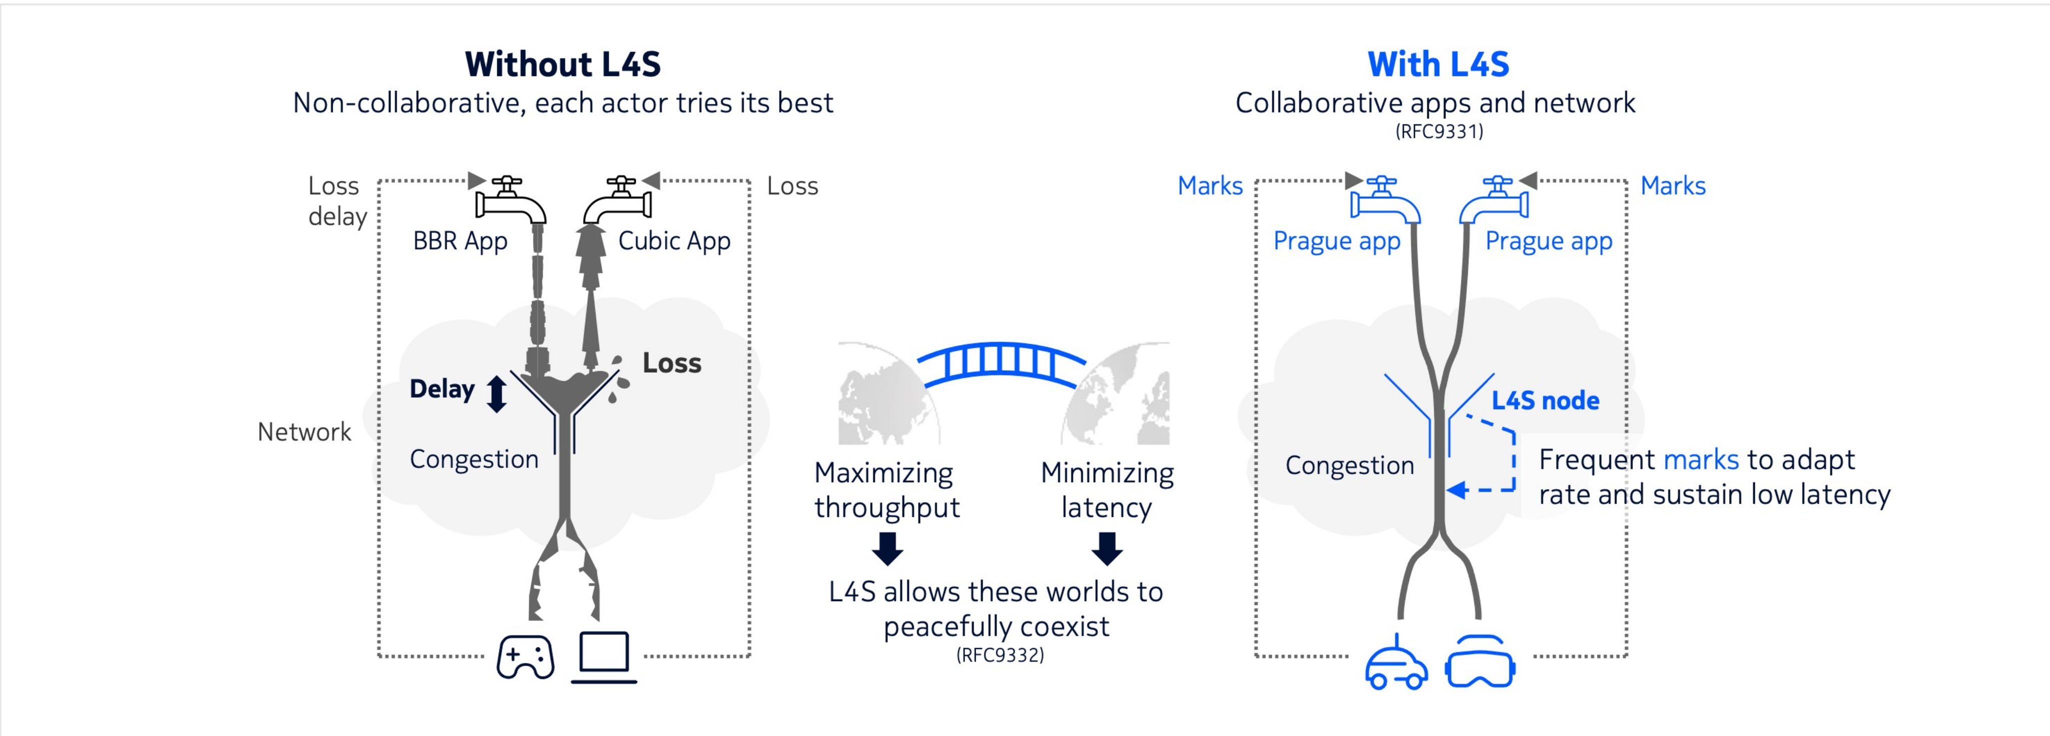

Dual Queue Basics

Leverage CE marks to provide senders with a faster congestion / delay signal than would be possible with packet drops today

(image from Bob Briscoe)

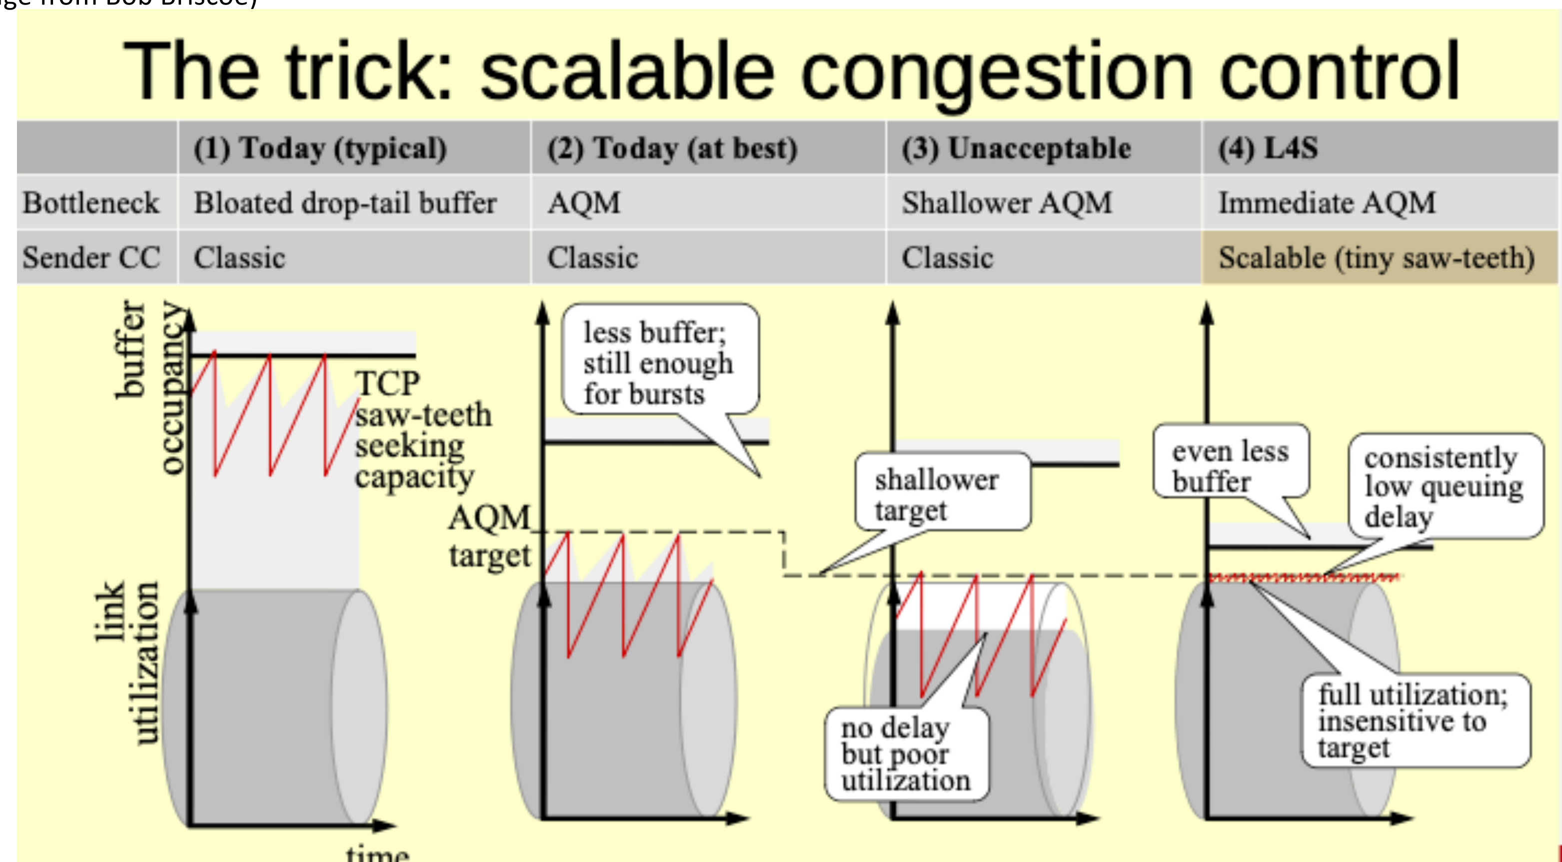

	(1) Today (typical)	(2) Today (at best)	(3) Unacceptable	(4) L4S
Bottleneck	Bloated drop-tail buffer	AQM	Shallower AQM	Immediate AQM
Sender CC	Classic	Classic	Classic	Scalable (tiny saw-teeth)

Two Options Low Latency Networking

1. Low Latency, Low Loss, Scalable Throughput (**L4S**)
    - Intended for: high bitrate latency sensitive flows
    - Server requirements: scalable congestion control algorithm, such as TCP-Prague, so the sender will act on Congestion Experienced (CE) marks

2. Non-Queue-Building Per Hop Behavior (**NQB**)
    - Intended for: low bitrate latency sensitive flows
    - Server requirements: None

References:
https://www.rfc-editor.org/rfc/rfc9330.html
https://www.rfc-editor.org/rfc/rfc9331.html
https://www.rfc-editor.org/rfc/rfc9332.html
https://datatracker.ietf.org/doc/draft-ietf-tsvwg-nqb/
https://github.com/jlivingood/IETF-L4S-Deployment/blob/main/App-Developer-Guide.md
https://github.com/jlivingood/IETF-L4S-Deployment/blob/main/Network-Config-Guide.md

Comcast Field Trial Observations

- No classic queue starvation issues observed (open to additional test suggestions)
- Working as expected so far – one app has identified excessive CE marking that will be fixed in an app update
- Cloud-native interactive applications seem to improve more than "legacy" apps
- More field trial work planned the next 90 days…

Next slides: some specific numbers…

Low Latency Networking Trial: Milestones

- 1st employee modem provisioned 7/11/23
- 1st customer modem 8/14/23
- 1st interdomain DSCP-45 packets 10/5/23 (with Valve)
- 2nd group of customers in Wave 2 added 10/11/23

Bringing total to ~200 customers

0xb7 = 101101 = NQB, DSCP-45 with CE mark

```
06:11:19.495711 IP (tos 0xb7,CE, ttl 108, id 64623, offset 0, flags [none], proto UDP (17), length 186)
    98.52.200.237.59119 > 205.196.6.213.27027: UDP, length 158
```

0xb5 = 10110101 = NQB, DSCP-45 with ECT(1) mark

```
06:11:19.639254 IP (tos 0xb5,ECT(1), ttl 108, id 64627, offset 0, flags [none], proto UDP (17), length 186)
    98.52.200.237.59119 > 205.196.6.213.27027: UDP, length 158
```

IMP US LLD Results --> 50%+ lag reduction

LLD Production Testing

About the Testing

This dashboard represents the data acquired from **Ultra-Low-Latency (ULL)** tests run on the list of employee and customer devices participating in the **Low-Latency-DOCSIS (LLD)** trial. Devices must be at a firmware version of at least 6.5 and either CGM4331COM (TCH XB7) or CGM4981COM (TCH XB8).

Each device will have six runs per day. One **run** consists of three different IMP measurements:

1. Latency-only DSCP 8
2. Latency-only DSCP 45
3. ULL upstream with DSCP 8 and 45

The latencies from the ULL measurements represent the Netperf metrics collected while consuming 100% of the device's upstream bandwidth. In contrast, the latency-only measurements do not congest the upstream bandwidth at all. We focus on the 99th percentile latency as that is the most accurate metric for determining customer experience.

About the Results

We look into standard bandwidth metrics such as PoA (percent of the advertised tier reached). Most importantly, we want to analyze the performances between queue (DSCP 8) and non-queue building (DSCP 45) upstream loaded latencies.

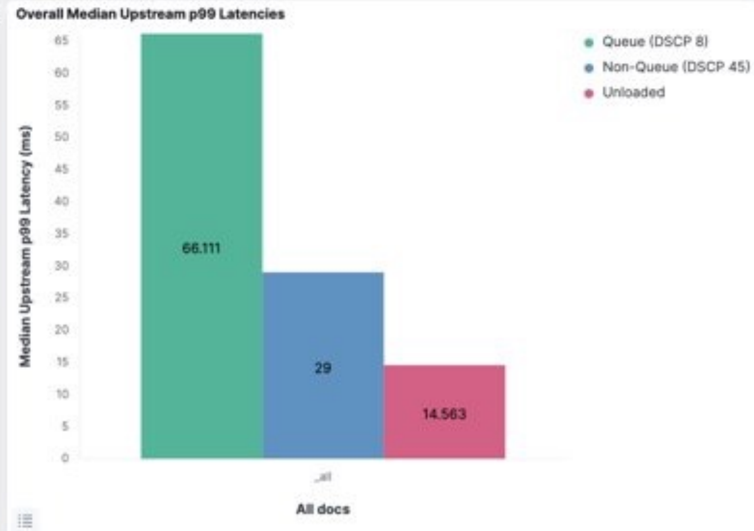

Unique Device Count - No Errors

30
CGM4331COM -

8
CGM4981COM -

Total ULL Test Count

4,752
Total Test Count - No Errors

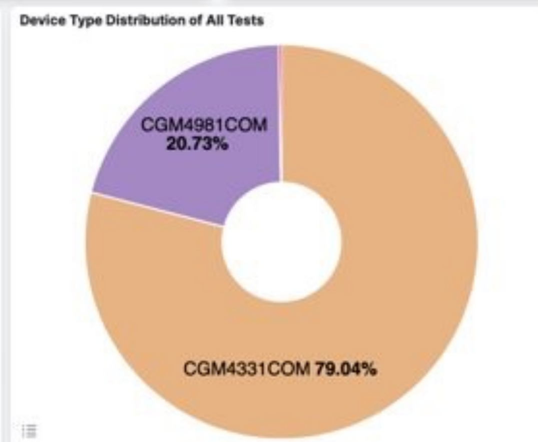

Lag Spikes
~20 ms vs 225 ms!

SUMMARY:

NVIDIA EXPERIMENTS WITH DIFFERENT MARKING

Low Latency Queue

- Significantly lower jitter – MUCH more consistent and reliable QoE
- Very low working latency for classic traffic AND low latency traffic
- Strong independent validation
- Makes cloud gaming truly viable

NVIDIA Working Latency

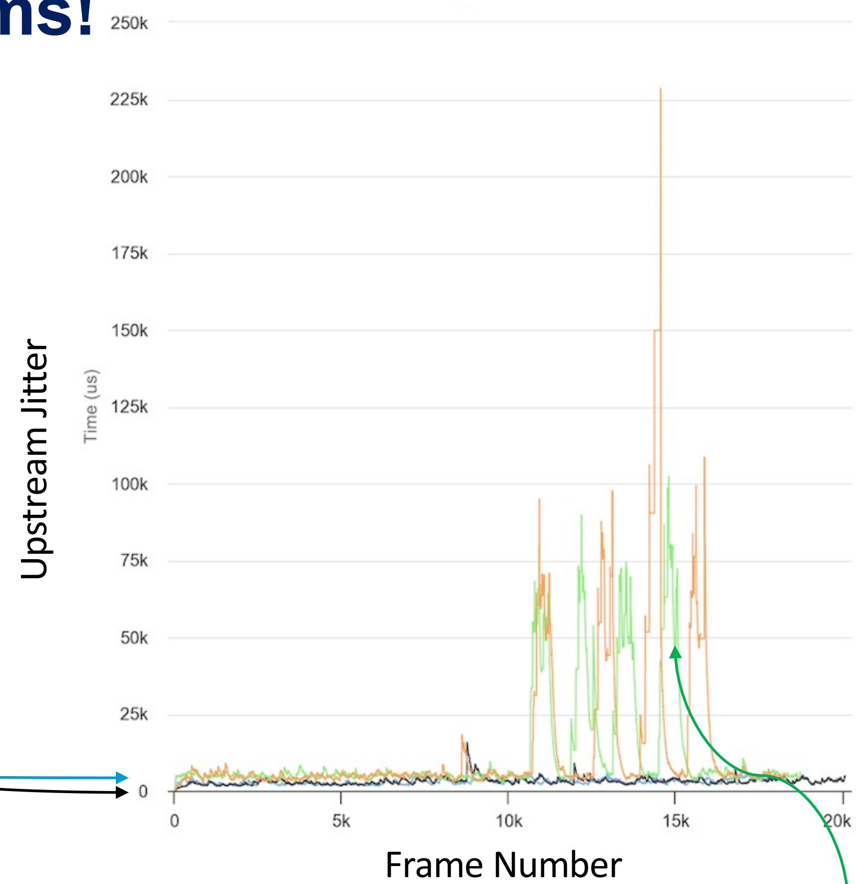

L4S ECT(1) flows show nominal latency increases

Flow loaded with classic queue traffic including several TCP file transfers with no L4S Congestion Experienced (CE) in DS vCMTS queues

Flow Marking
- Orange = Downstream no marking – normal working conditions
- Green = Upstream no marking – normal working conditions
- Black & Blue = DS & US with LLD marking – normal working conditions

NVIDIA GEFORCE NOW

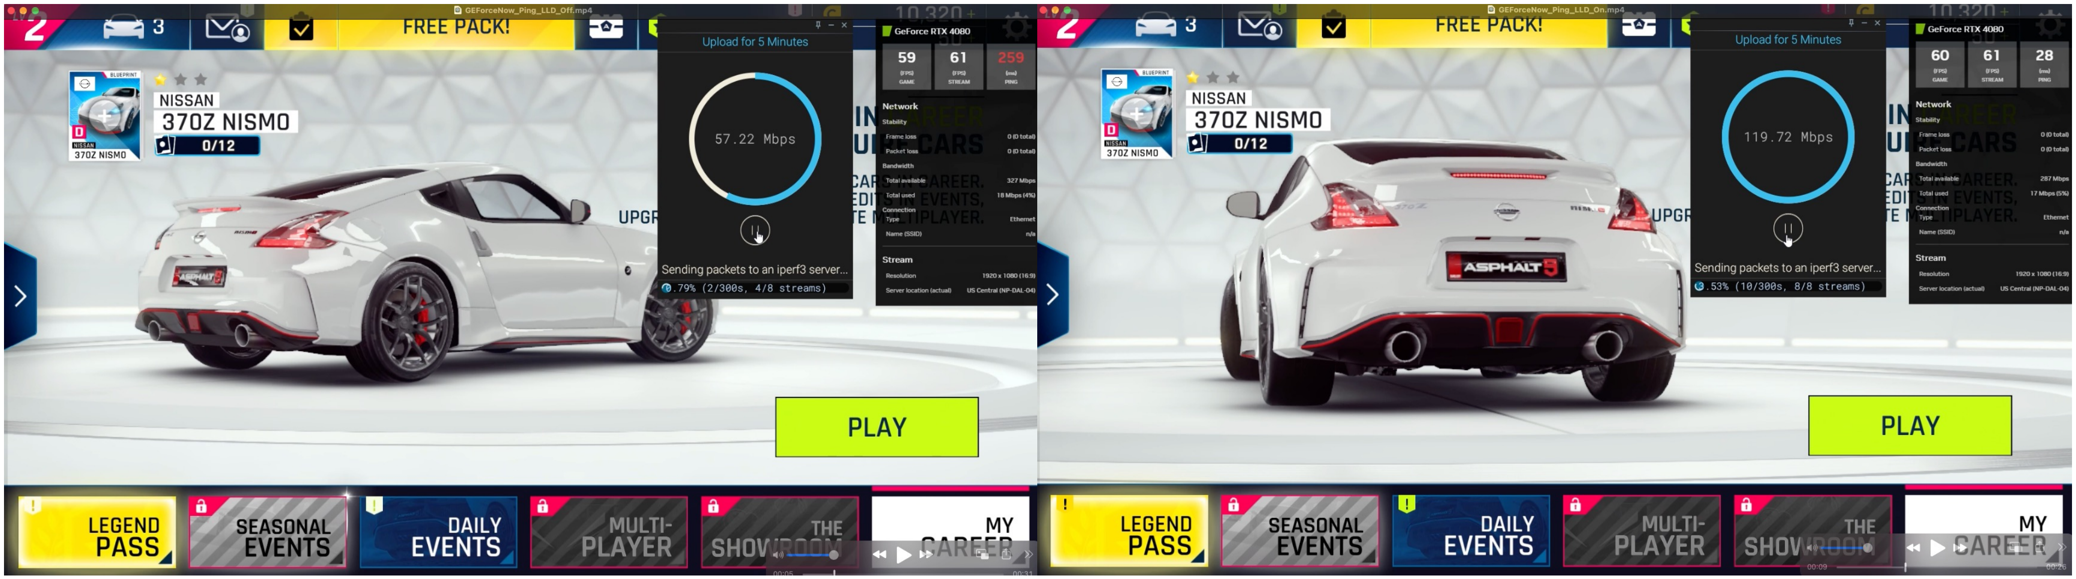

Example cloud game:

- Upstream traffic generator – creates bursty, variable upstream flows
- Left is without L4S – ping spiking to 259 ms, music & video out of sync and generally buffering
- Right is with L4S – stable ping and seamless QoE

Valve/Steam *Counterstrike* Results – Down to Roughly Idle!

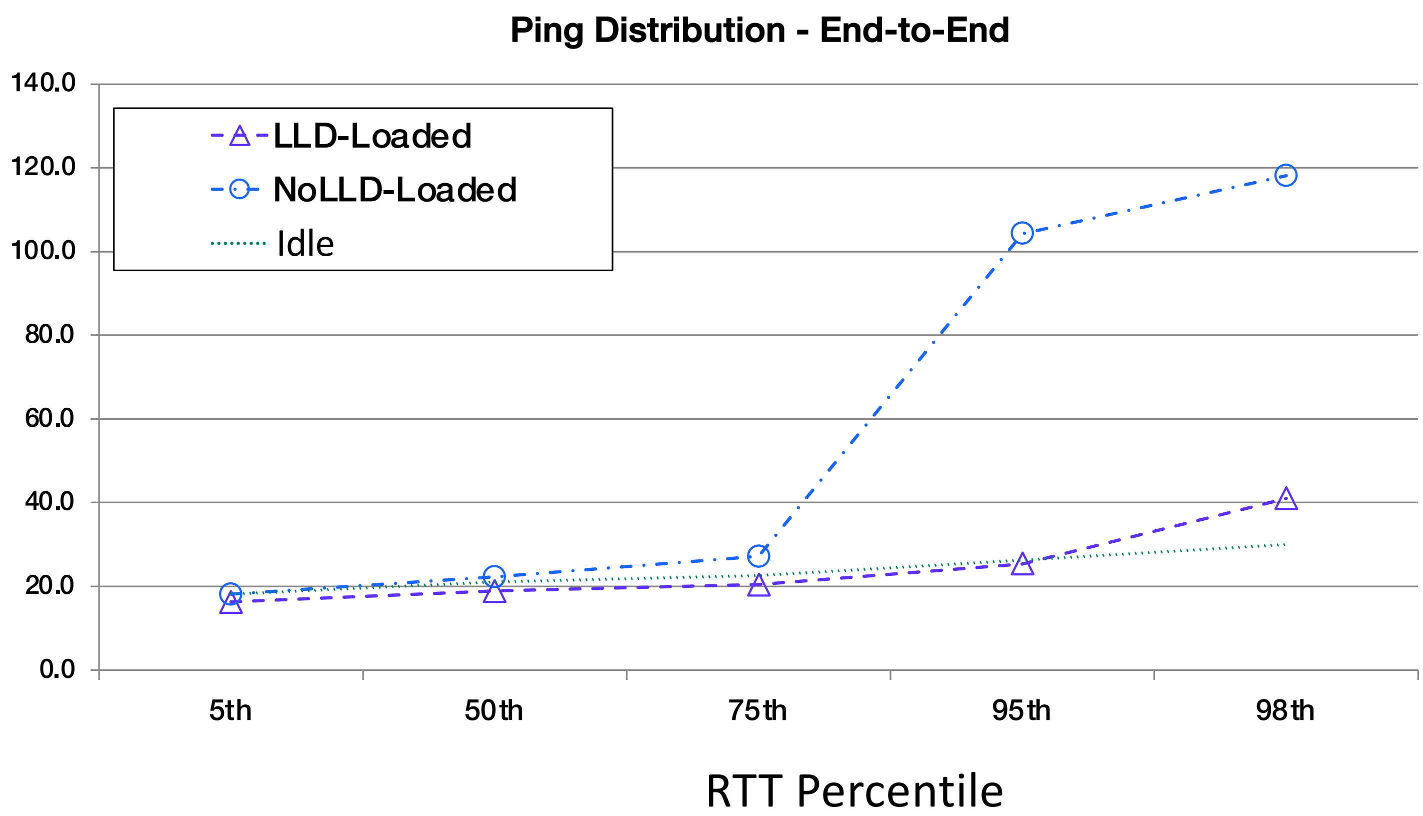

Conclusions & Next Steps

- Comcast to continue field trails through March 2024
- Preparing to operationally scale to **millions** of users
- App provider interest is high – expect key platforms to adopt
- **Network Operators**: The time is NOW to start working on dual queue low latency networking. You need to take steps to help deployment.

  1. Validate that the ECN header is not modified as packets traverse your network – that will enable L4S to work e2e
  2. Update peering edge router policies to allow DSCP-45 packet marking to pass into and out of your network – that will enable NQB to work e2e – and be sure that DSCP-45 is treated as Best Effort priority (same as other internet traffic)