IRTF - MAPRG

Comcast
Dual Queue Low Latency
Field Trial Measurements

IETF 119 – March 2024

Presented by Jason Livingood

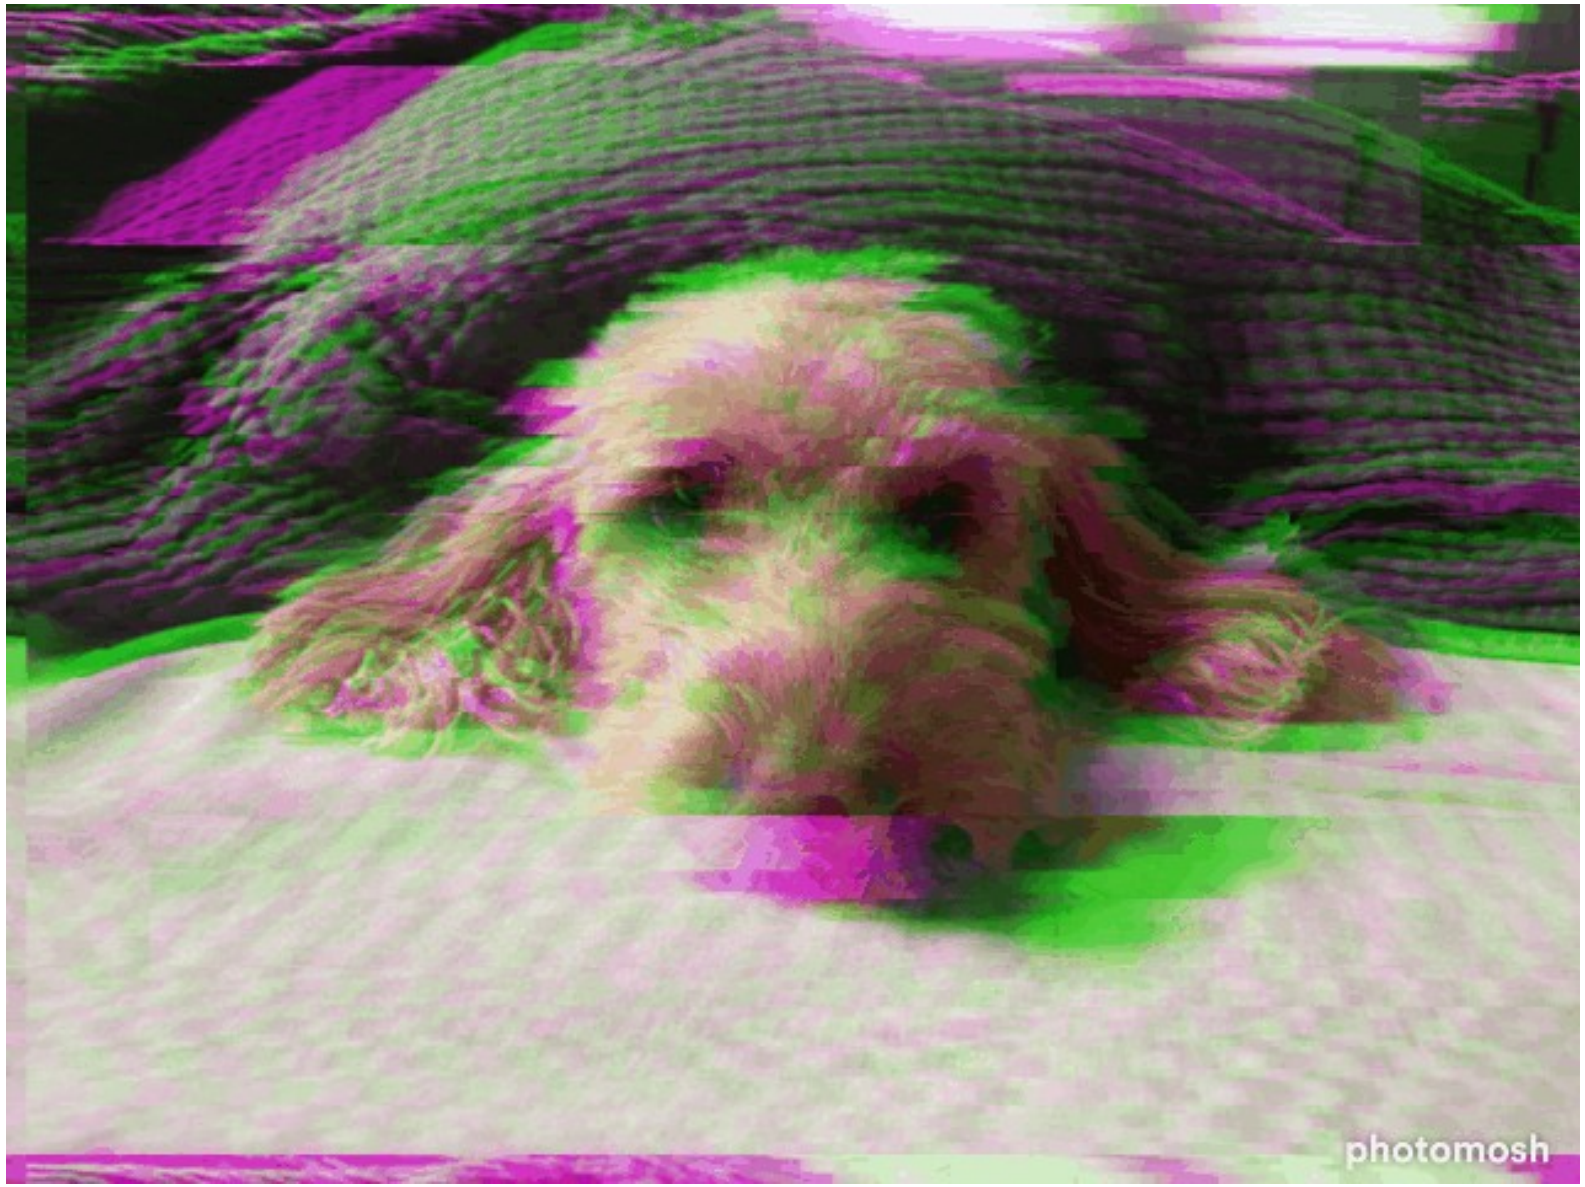

COMCAST

Background: L4S & NQB

L4S: See RFC 9330, 9331, 9332

NQB: See draft-ietf-tsvwg-nqb-22

In short:

- Focused on latency-sensitive traffic, as signaled by end user applications
- Creates a new 2nd network queue, to be used at bottlenecks (i.e. home CPE)
- LL queue is same priority (best effort) and shares same bandwidth
- L4S can handle high throughput demands (i.e. Facetime, NVIDIA GeForce NOW) – but requires the server to have a suitable congestion controller
- NQB is intended for low bitrate flows (i.e. DNS, signaling traffic)
- Marking done by the application
  - L4S = ECN header, ECT(1) and CE
  - NQB = DSCP header, DSCP-45

Comcast's Low Latency Field Trial

Timeline:

- Announced in June 2023 – high customer interest
- Phase 1: July - October 2023 – Upstream, WiFi LAN (AC_VI), DSCP-45 PoC
- Phase 2: November 2023 – Today – Downstream, Automation, Full DSCP-45

Technical Requirements:

- User connected to virtual CMTS platform (vCMTS)
- Compatible cable modem, both ISP-provided and COAM
- Can be located anywhere in the network geographically

Customers Assigned Structured Weekly Activities:

1. Some tests auto-report, others required customers submit results in a form.
2. We also asked trial partners for app-layer QoE stats.
3. We also deployed 3rd party test probes

COMCAST

Approach to Test & Measurement

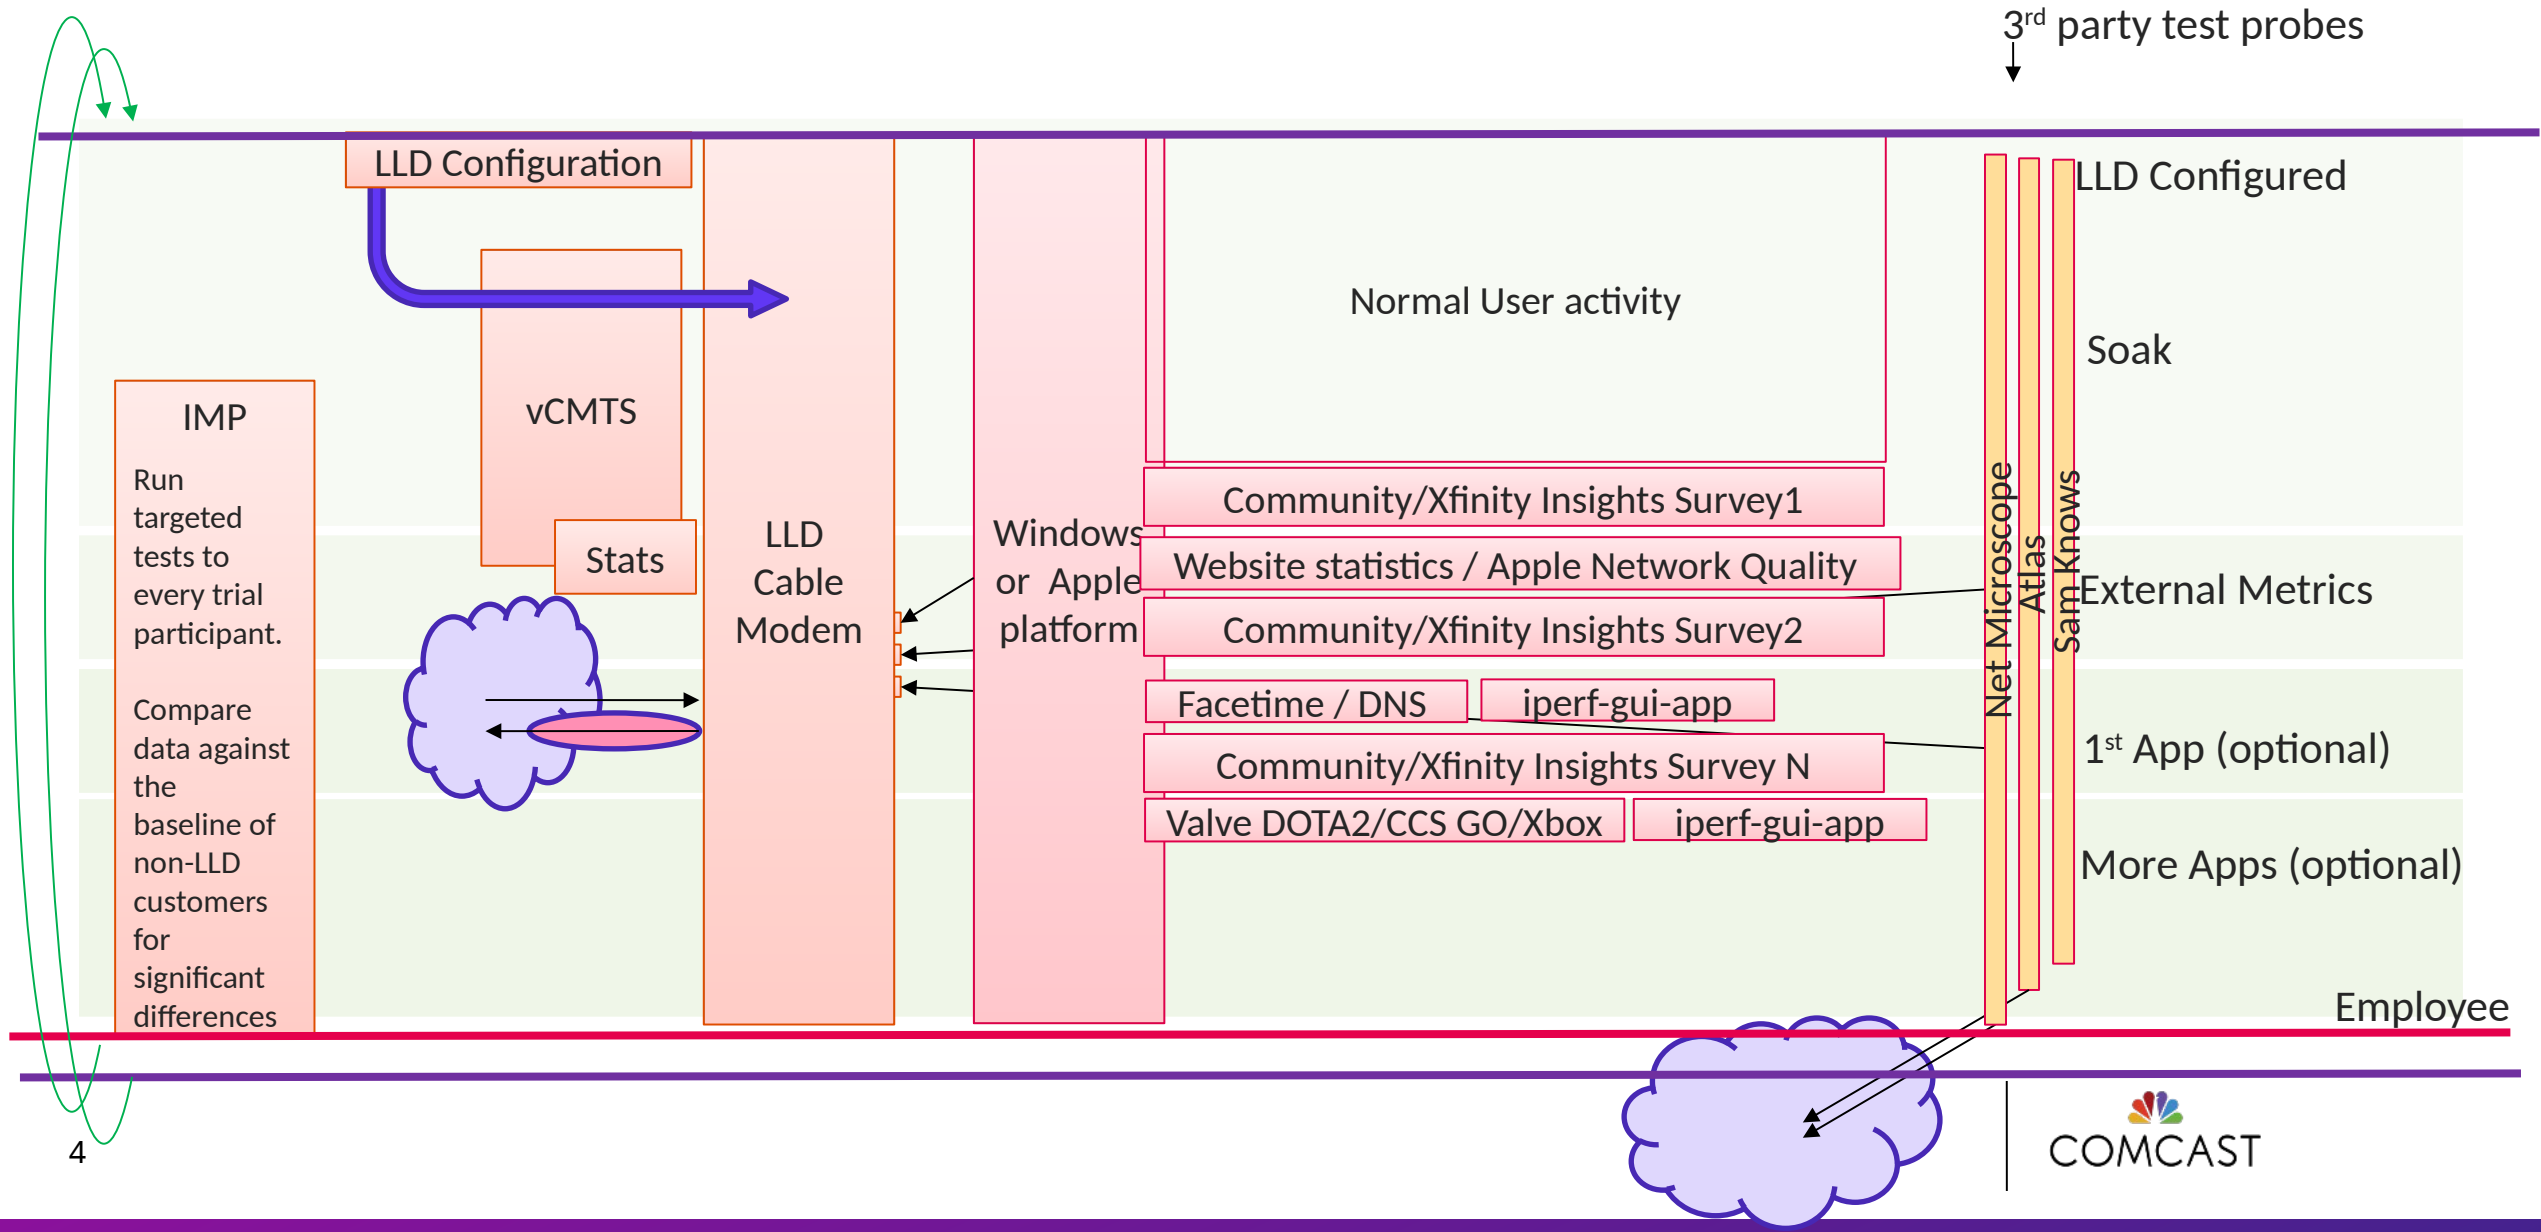

COMCAST

Approach to Test & Measurement – Another View

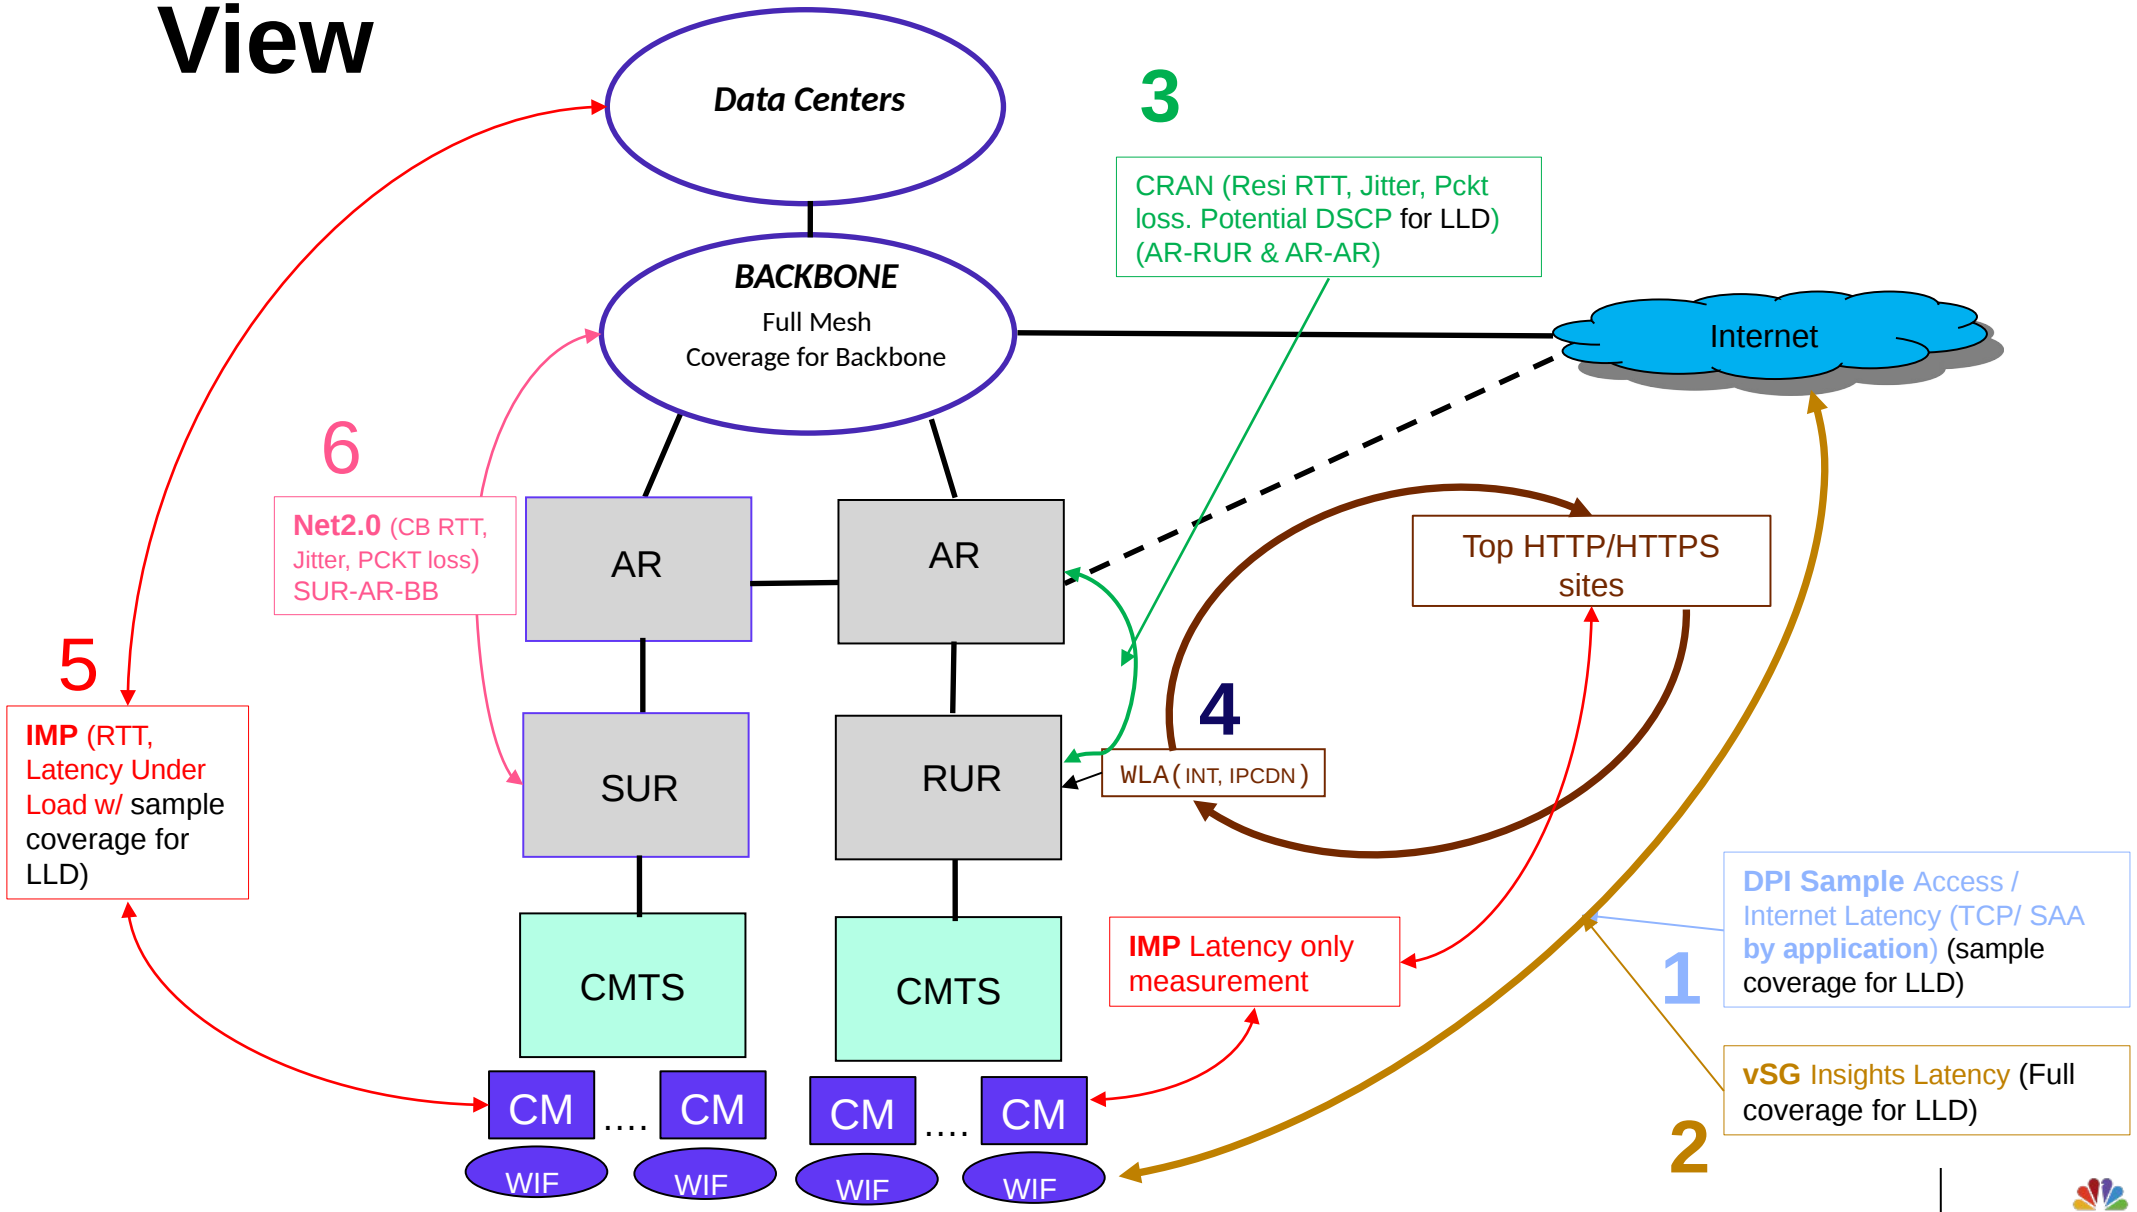

COMCAST

Selected Measurement Findings

Discovery of a Measurement Design Error:

- Testing ECN gave us the impression all was well (no bleaching) but our test was uni-directionally (upstream).
- We tested bi-directional and found we were actually bleaching downstream marks, so we had to fix that on all CMTS Service Class Names (SCNs). At the same time, we also fixed a longstanding DSCP leak (CS1->WMM AC_BK) that we had just learned about & would have posed problems in testing LL. Every cable modem in the network was updated.

Confirmed no classic queue starvation (key question) – mainly via PCAPs and pulling interface stats.

Initially saw some LL queue starvation in the measurement data – prompted a major update of the cable modem configuration file

Unlike AQM (DOCSIS-PIE) – LLD is not on/off:

```
"direction": "upstream",
"max-concatenated-burst": 12345,
"max-traffic-burst": 12345,
"max-traffic-rate": 12345000,
"maximum-buffer": X,
"min-reserved-packet": X,
"multiplier-bytes-requested": X,
"multiplier-contention-request-window": X,
"peak-traffic-rate": X,
"priority": X,
"service-class-name": ”ABCD",
"sf-aqm-latency-target": XX,
"tos-and-mask": ”XX",
"tos-or-mask": ”XX",
```

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

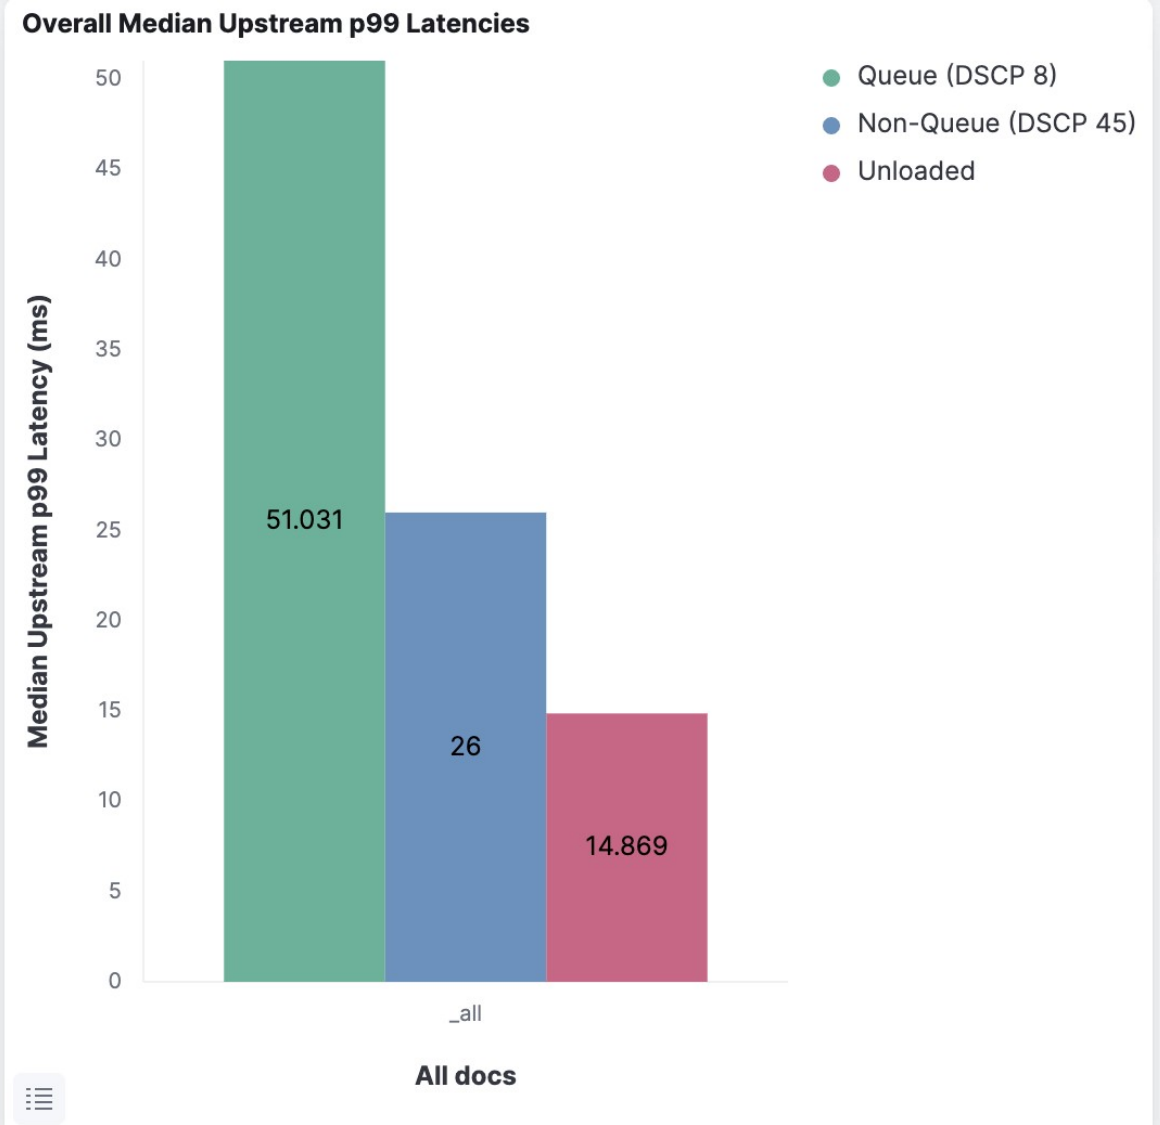

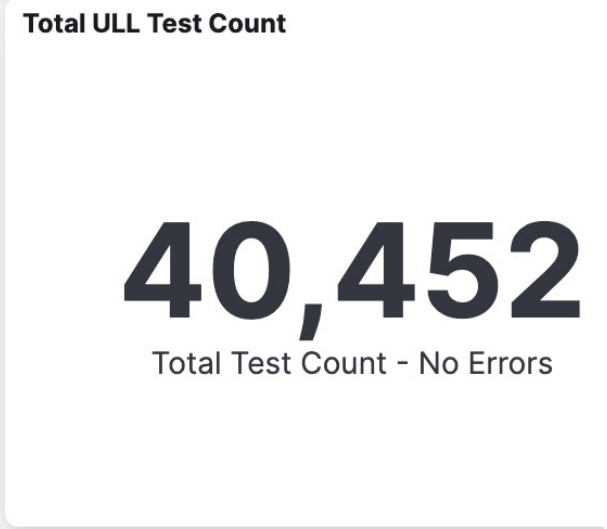

Ignore the DSCP notation – this is not limited to NQB – the “DSCP 45” label represents L4S and NQB.

DSCP-8 is how best effort residential internet traffic is marked in the network (like every network, we have an internal DSCP code point usage scheme).

These synthetic tests are an important tool to validate latency performance. But the best data come from applications themselves.

COMCAST

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

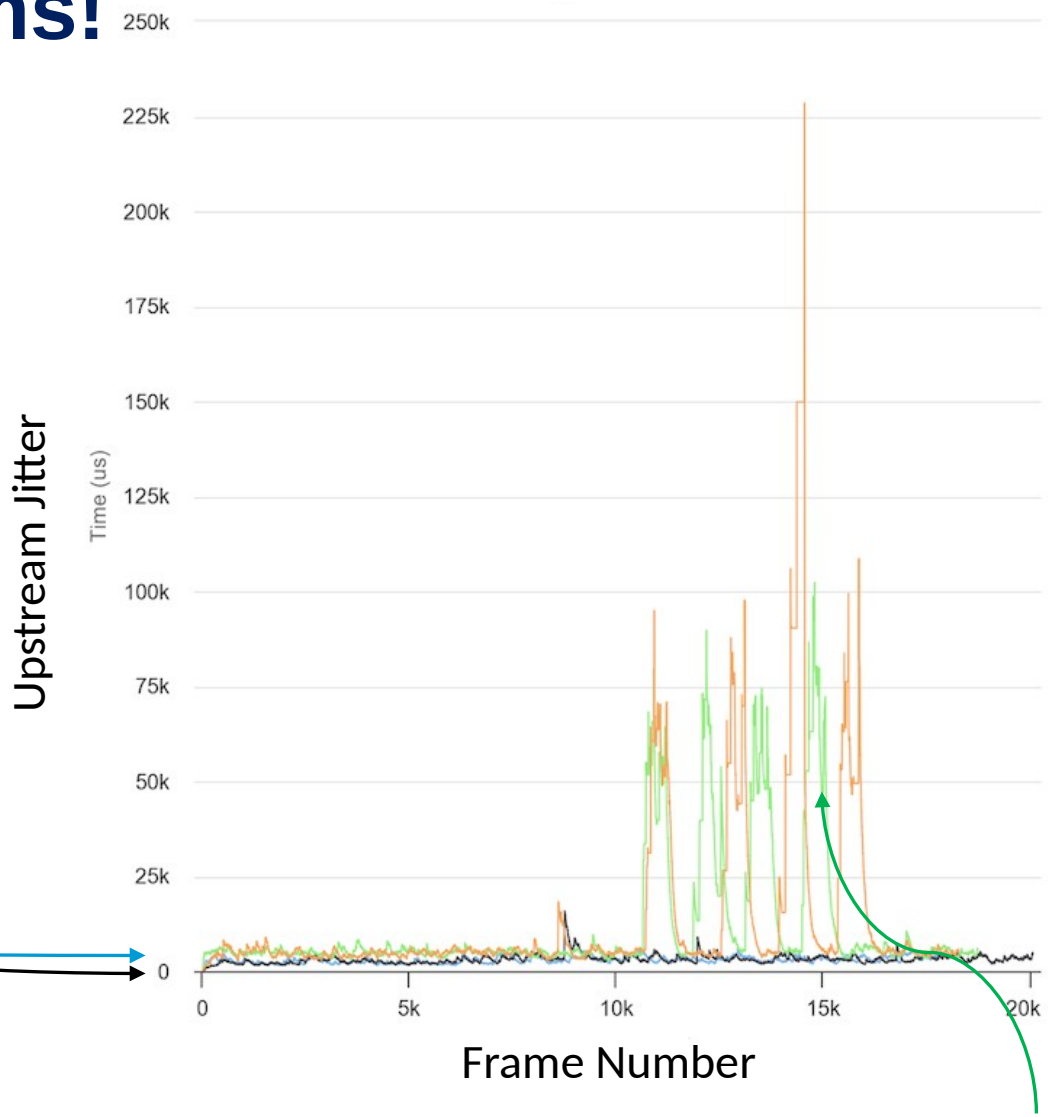

L4S ECT(1) flows show nominal latency increases

Flow 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

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

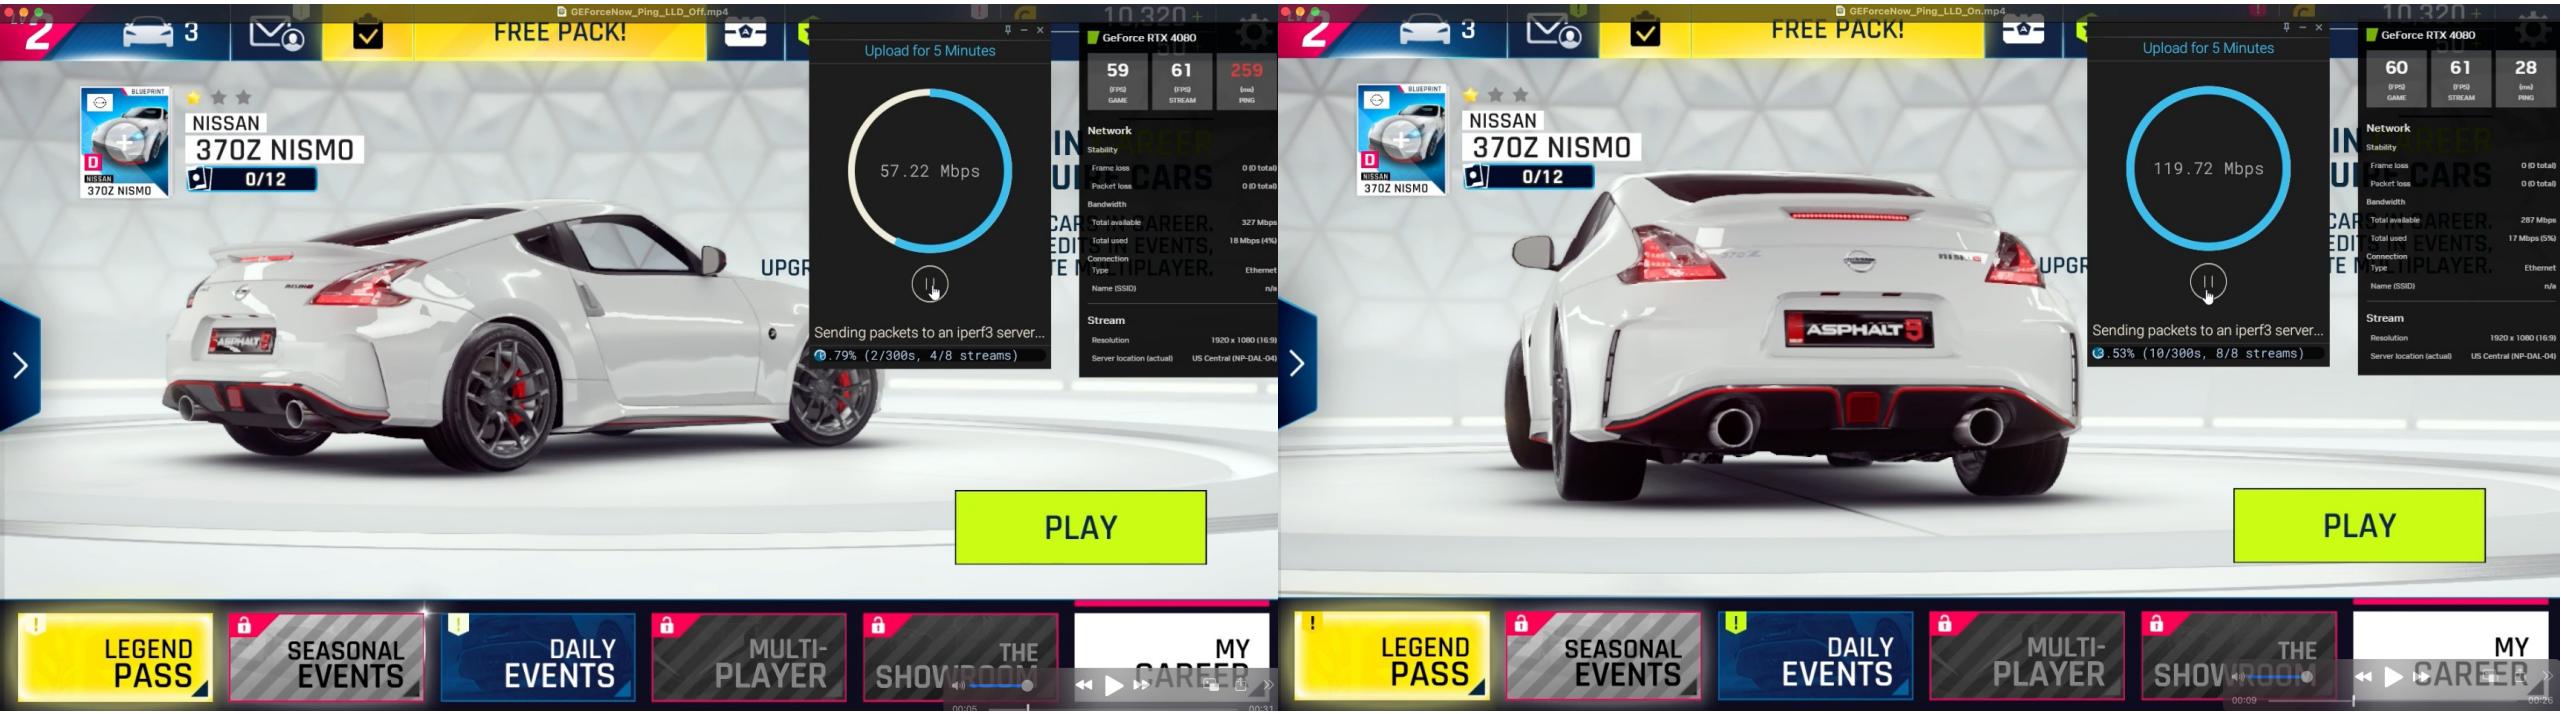

Example cloud game:
- Upstream traffic generator – creates bursty, variable upstream flows
- Left is without L4S – ping spiking to 259 ms, music & video out of sync and generally buffering
- Right is with L4S – stable ping and seamless QoE

COMCAST

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

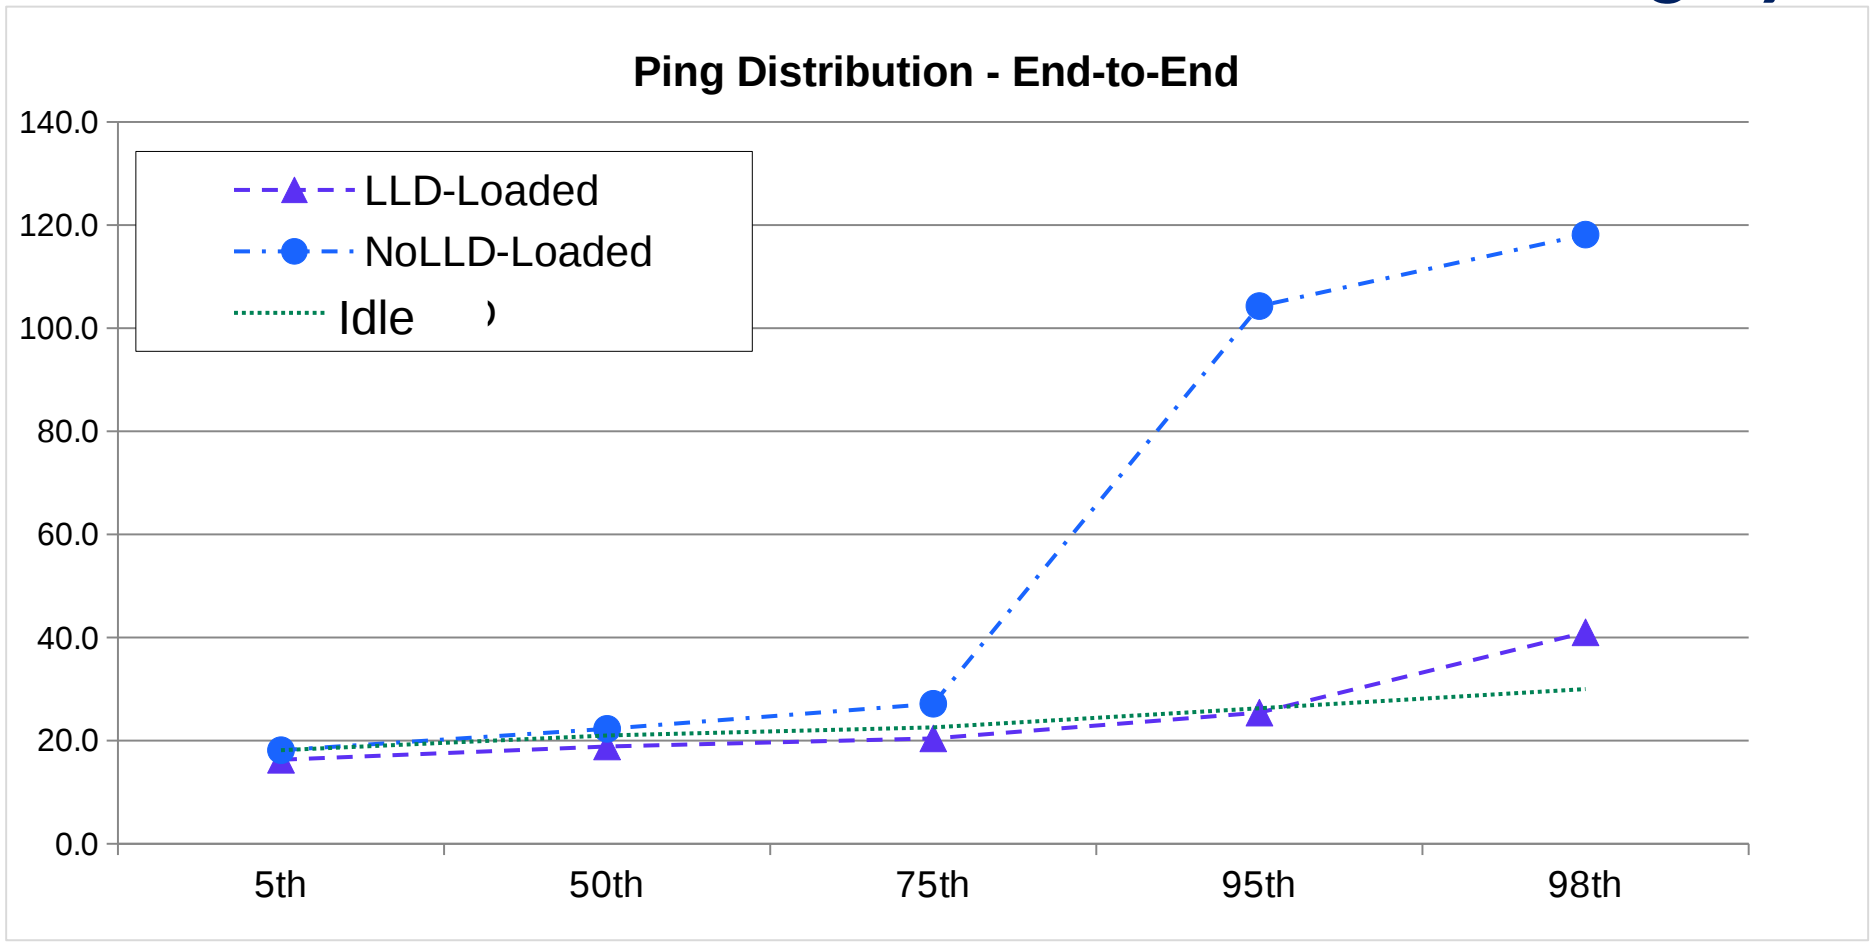

COMCAST

Conclusions & Next Steps

- Testing to continue for a TBD period of time
- Lots of measurements. New bootfile deployment = re-running all prior tests
- Preparing to operationally scale to **millions** of users