L4S

Simple & scalable E2E support of low-latency traffic

Koen De Schepper, Research Lead

Werner Coomans, Technology Advisor

May 21st, 2024 – RIPE88

https://bell-labs.com/l4s

L4S = Low Latency, Low Loss, and Scalable throughput

A new IETF internet protocol to reduce queuing delay to near-zero values

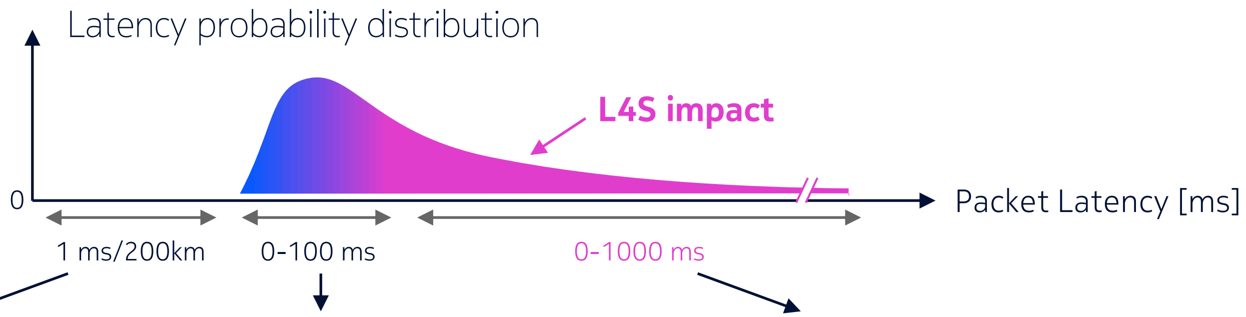

Propagation delay

- Limited by speed of light
- Addressed by move to **edge clouds**

Interface delay

- Limited by PHY & MAC-layer implementations
- Addressed by new **communication technology standards** (e.g., 5G, PON)

Queuing delay

- Biggest source of **latency variations**, caused by queuing in network buffers
- Addressable by avoiding (GBR/slicing) or managing queuing delay (AQM, **IETF L4S**)

NOKIA BELL LABS

The goal of L4S is to reduce “working latency”

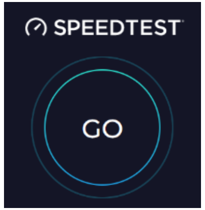

&

ping -t 8.8.8.8

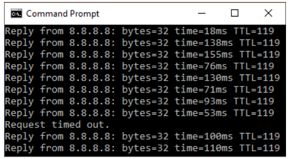
Command Prompt

```
Reply from 8.8.8.8: bytes=32 time=18ms TTL=119
Reply from 8.8.8.8: bytes=32 time=138ms TTL=119
Reply from 8.8.8.8: bytes=32 time=155ms TTL=119
Reply from 8.8.8.8: bytes=32 time=76ms TTL=119
Reply from 8.8.8.8: bytes=32 time=130ms TTL=119
Reply from 8.8.8.8: bytes=32 time=71ms TTL=119
Reply from 8.8.8.8: bytes=32 time=93ms TTL=119
Reply from 8.8.8.8: bytes=32 time=53ms TTL=119
Request timed out.
Reply from 8.8.8.8: bytes=32 time=100ms TTL=119
Reply from 8.8.8.8: bytes=32 time=110ms TTL=119
```

Ping latency during Speedtest

on a commercial 4G network

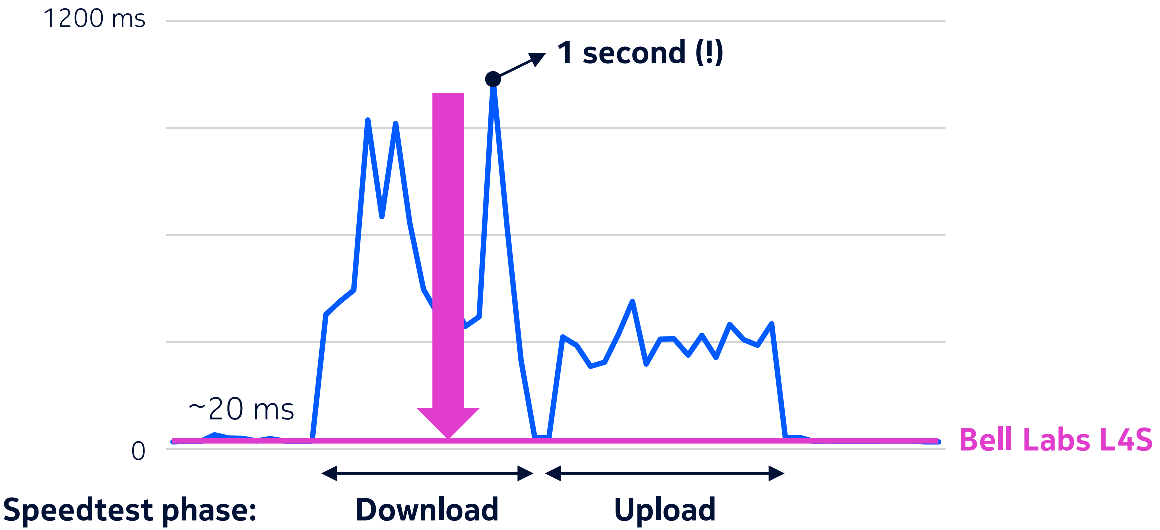

L4S can drastically improve the Quality-of-Experience

Of any application benefiting from a consistently low latency

As demonstrated by Apple:

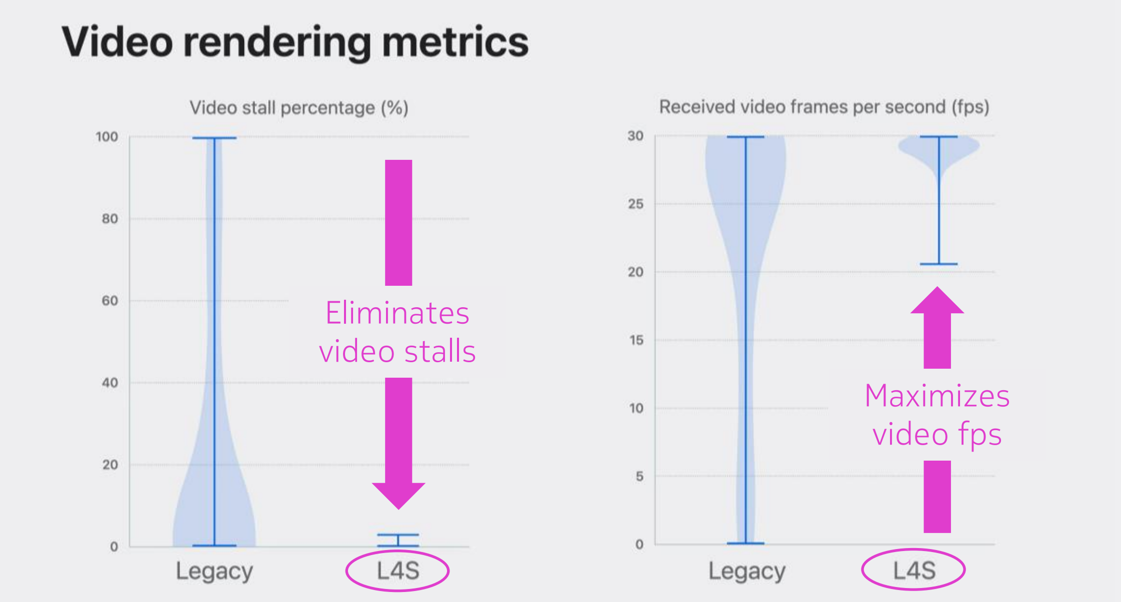

From: https://developer.apple.com/videos/play/wwdc2023/10004/

Application/OS players active in IETF L4S interops:

Apple, Nvidia, Meta, Google, Netflix, ...

NOKIA BELL LABS

L4S combines a new rate-adaptation algorithm in the application with ECN-marking-based network rate control

Application

New rate-adaptation algorithm

- “Prague” requirements [RFC 9331]
- Fine-grained rate-control from 100 kbps to infinity (even for very low RTTs)

Network

Immediate marking-based rate control

- No packet-drops, but marks
- Instant congestion-signaling from the network (no smoothing)
- No need to build a queue before marking

Coexisting and compatible with non-L4S traffic

- No starvation of non-L4S traffic
- Using marks to control L4S rate
- For example, using a dual-queue PI2 AQM [RFC 9332]

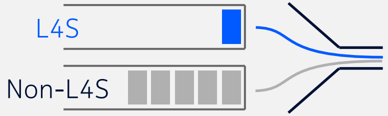

Classic AQMs & congestion-control face unavoidable trade-offs

Requiring a queue to limit rate variations, control rates, and limit packet loss

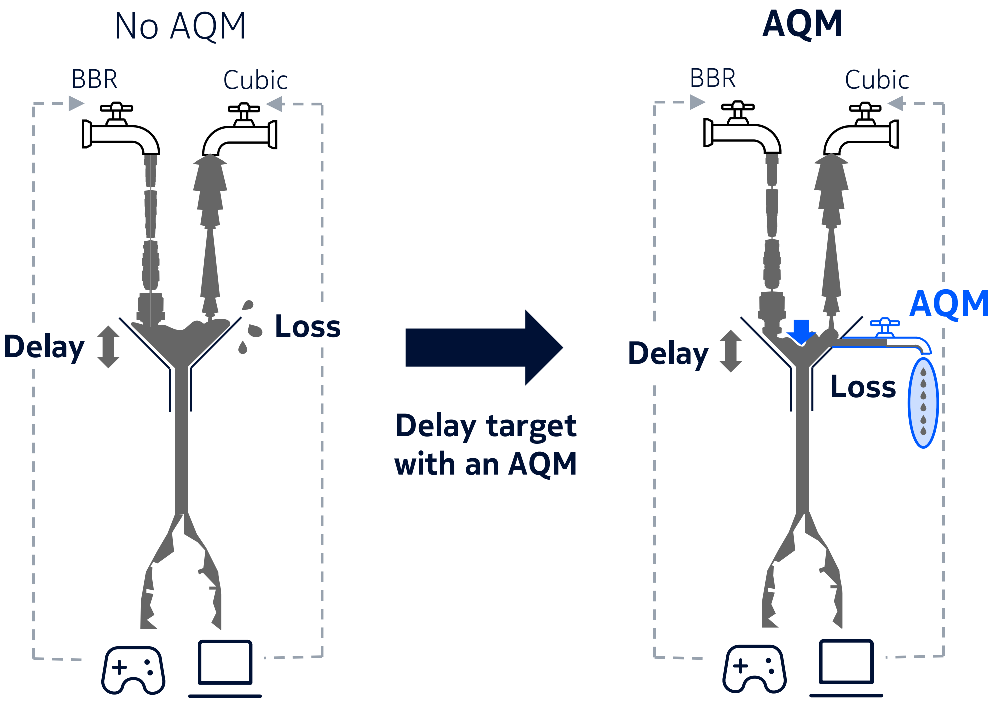

A queue is still needed to:

- Cover for data rate variations
- Control the rate of delay-based congestion-control algorithms
- Improve rate/RTT-fairness (bigger queue = more fairness)
- **Limit packet loss rate (lower latency requires higher loss)**

Using classic ECN instead of drops lowers the packet loss

But does not reduce latency

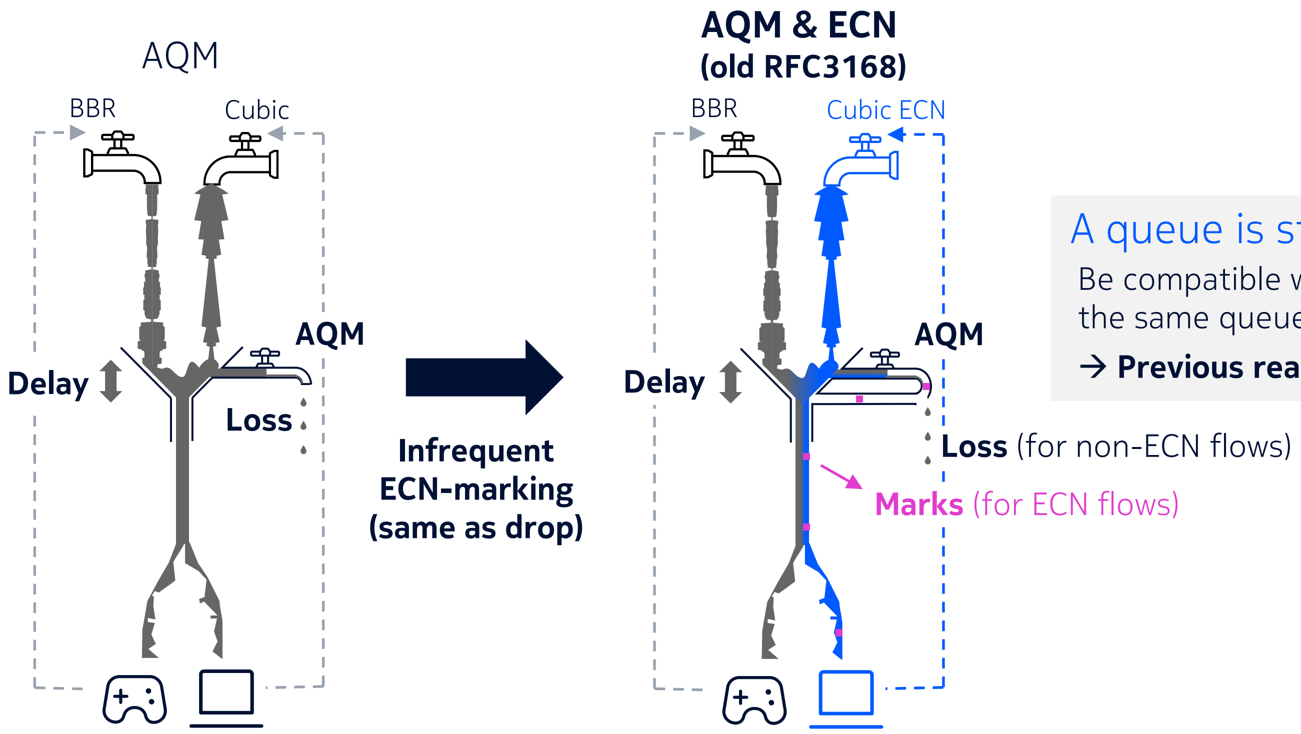

A queue is still needed to:

Be compatible with loss-based flows in the same queue

→ **Previous reasons still apply!**

L4S was inspired by Data Center TCP

But it required many changes to make it work on the public internet

Immediate AQM & ECN

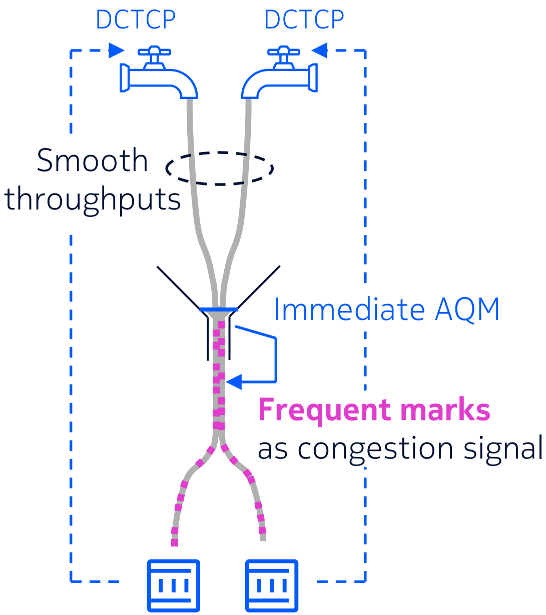

Data Center TCP

Enables low latency & smooth, high throughputs, BUT:

- **Cannot coexist with non-DCTCP traffic**
- **Doesn't work for (lower) internet rates and (higher) internet RTTs**

L4S solves this by introducing:

- **Coexistence and compatibility mechanisms** with non-L4S traffic (e.g., Dual PI2 [RFC9332])
- **Prague congestion-control**, adding e.g.:
  - Source-pacing
  - Burst-size limits
  - RTT-independence

L4S leverages a near-constant rate of congestion signals

Using rate-adaptation operating according to the Prague requirements

Adaptable across
a near-infinite range

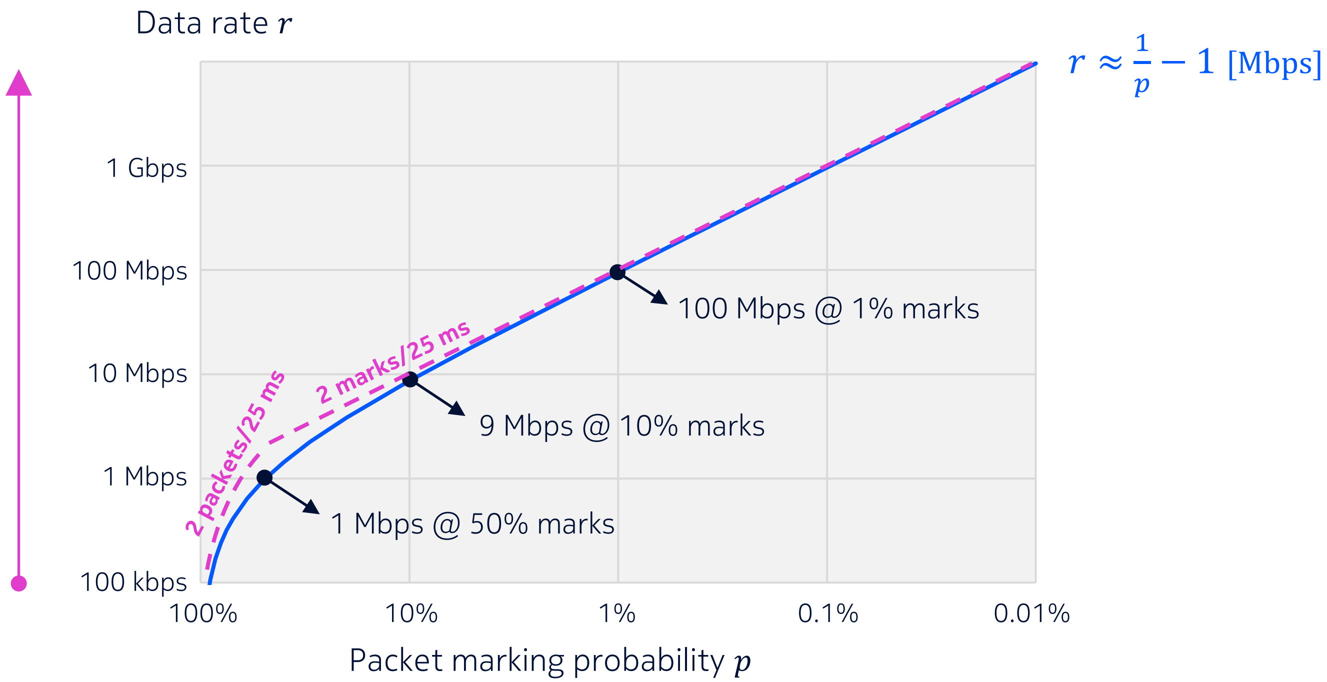

NOKIA BELL LABS

L4S is technology-agnostic

L4S packets are identified by the ECN bits in the IP header

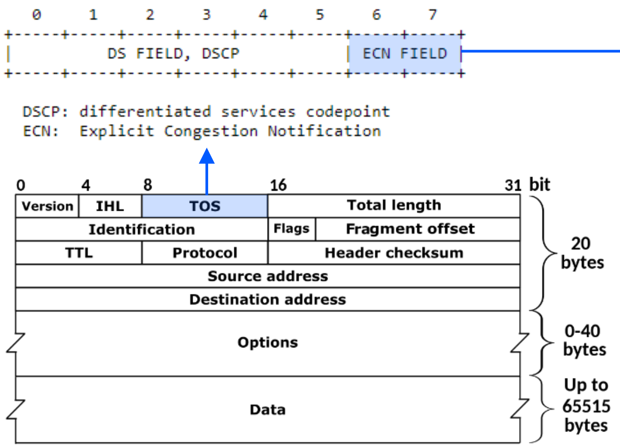

IPv4 packet structure

ECN field Value	Codepoint name	
00	Non-ECT	Not related to L4S
10	ECT(0)	Not related to L4S
01	**ECT(1)**	Set by source to identify as L4S packet
11	**CE**	Set by the network to signal congestion (through **marking** instead of dropping)

L4S offers a uniform rate-adaptation mechanism for applications

What does it mean to support L4S in a network node?

1. Removing "in-flight" jitter through **Prioritization**

2. Network feedback for source rate-adaptation

Classification

Identify L4S traffic

Isolation

Use separate queues for L4S packets and non-L4S packets

Prioritization

Allow the few L4S packets to skip ahead of the non-L4S packets

ECN marking

Proper L4S-ECN marking for coping with rate bottleneck(s)

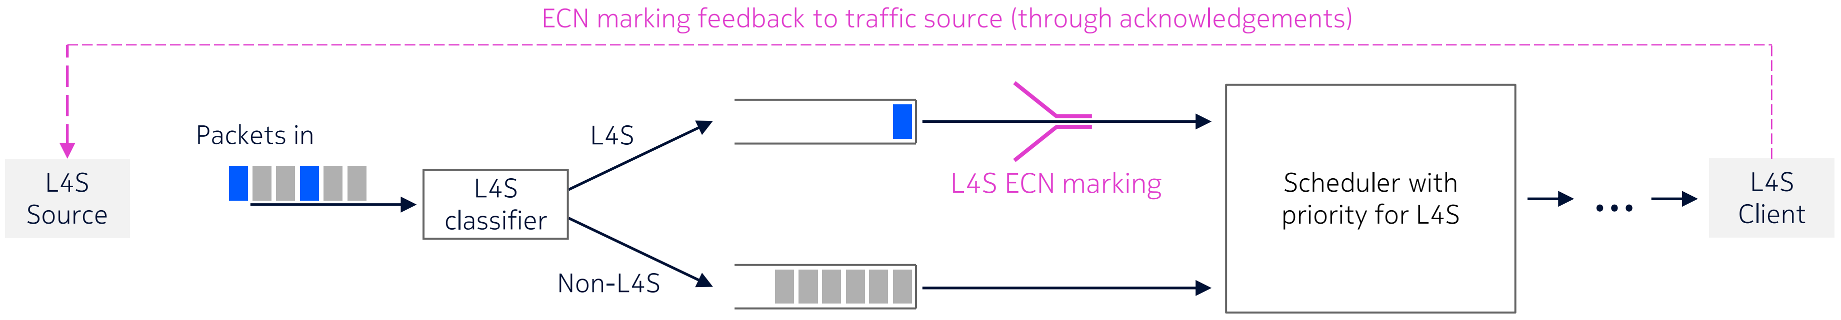

L4S enables scalability in low-latency service offerings

Through fast rate-adaptation, while safeguarding QoE

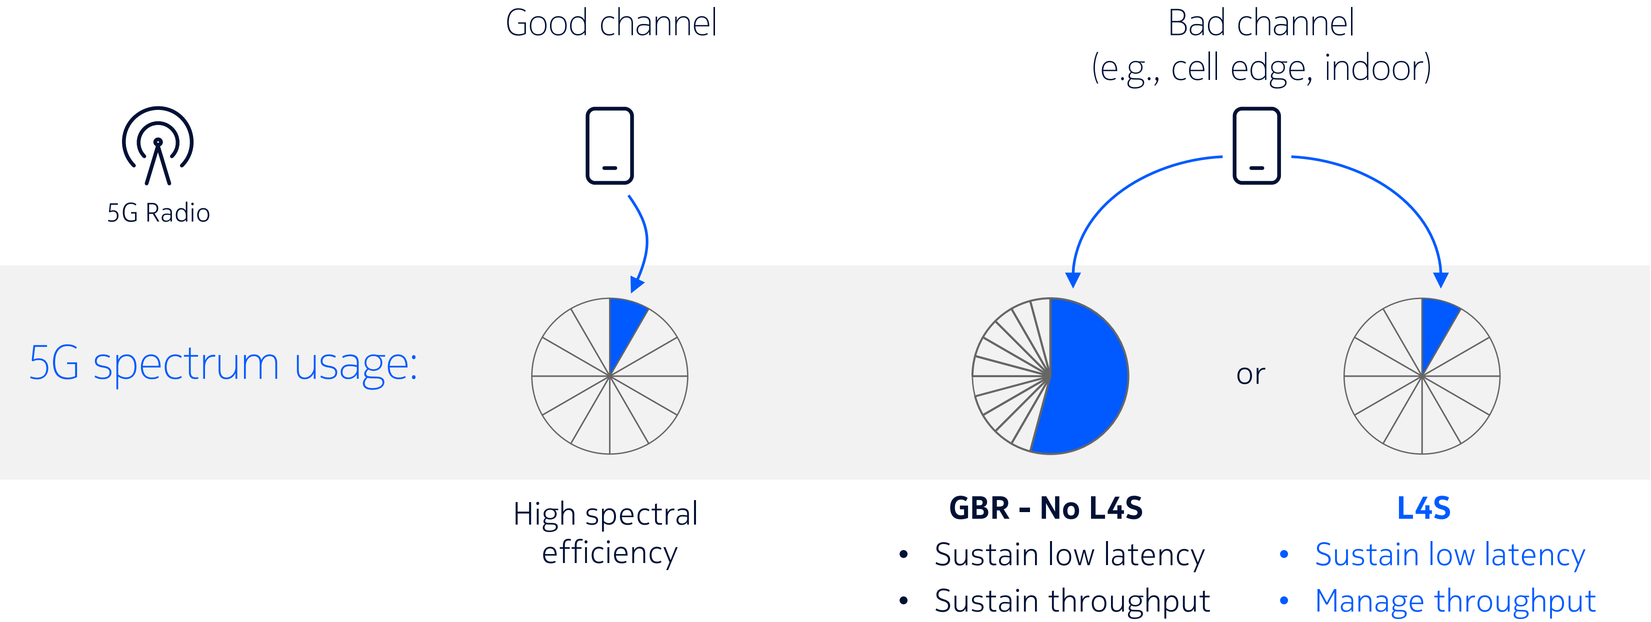

L4S enables scalability in low-latency service offerings

L4S flows will each get a flow-fair share of a common bottleneck

2 L4S traffic flows

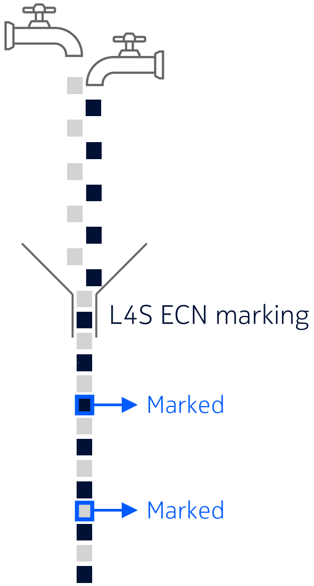

2x number of L4S flows

2x number of marks

½ throughput/flow

L4S ECN marking will automatically converge and result in **flow-fair capacity-sharing**

4 L4S traffic flows

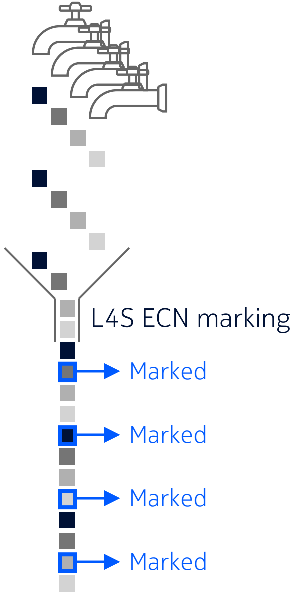

L4S support is not required on the entire end-to-end path

L4S support in access and in-home networks will yield biggest gains

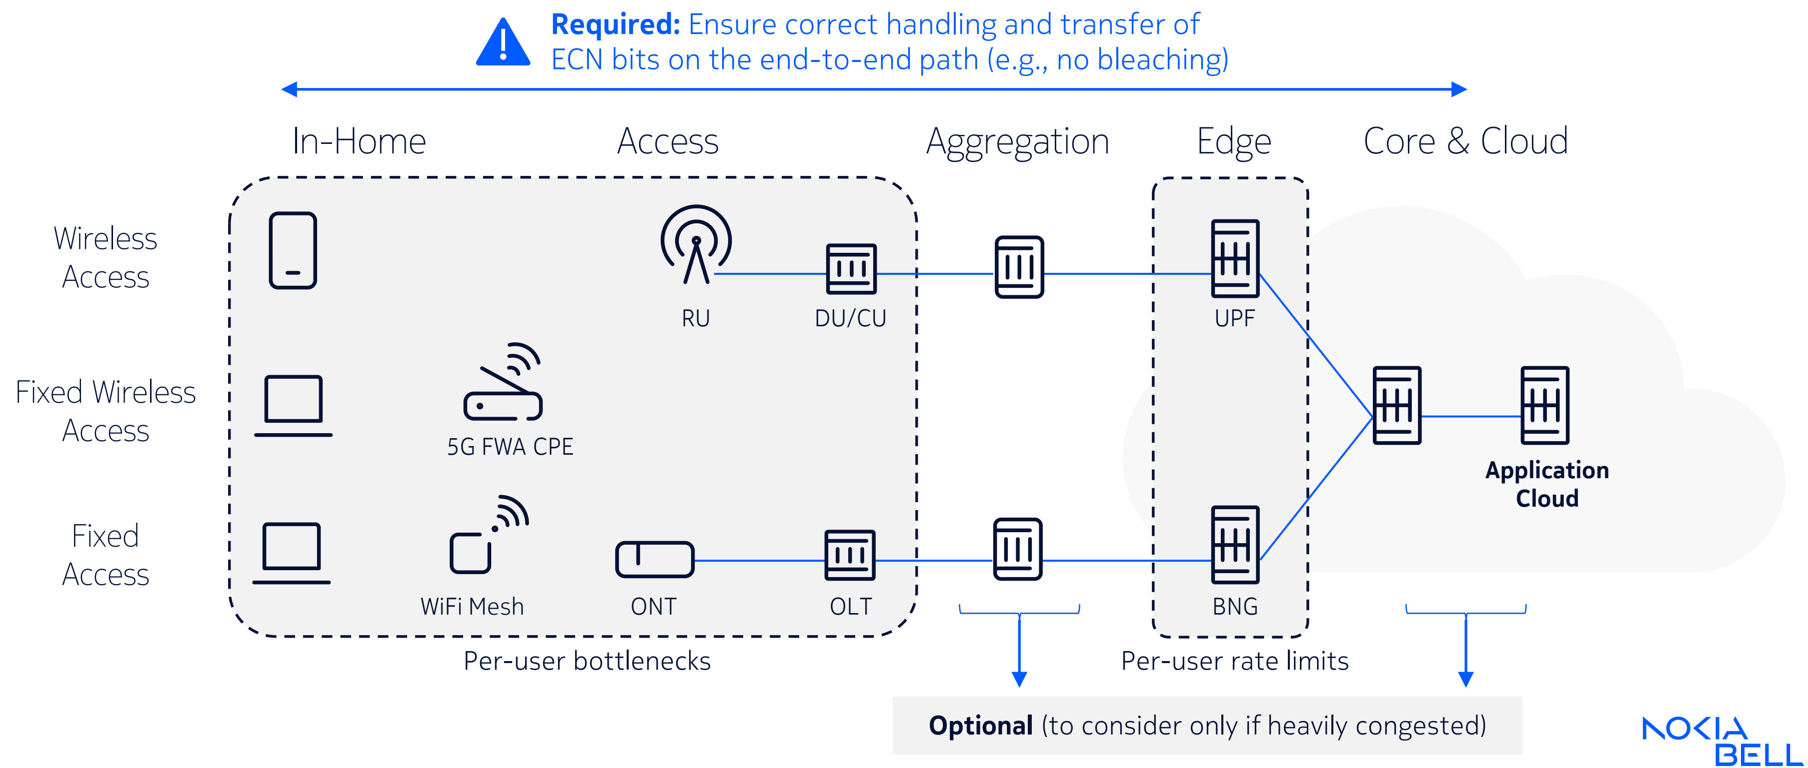

Adequate support of L4S in a network segment depends on the significance of the jitter it introduces

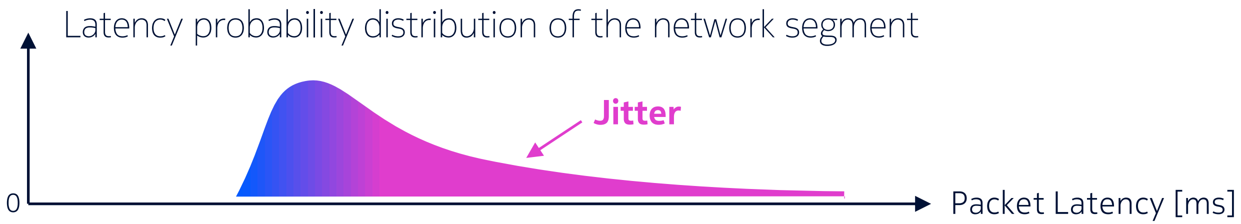

Jitter =	Insignificant	Significant, but transient	Significant and consistent
What to do	Nothing	Prioritize L4S	• Prioritize L4S • Support L4S-ECN marking – Signal bottleneck rate to source – And/or avoid starvation of lower-priority traffic

NOKIA BELL LABS

L4S support in a Layer-2 network

Can be done by ECN bit remapping at the network edges

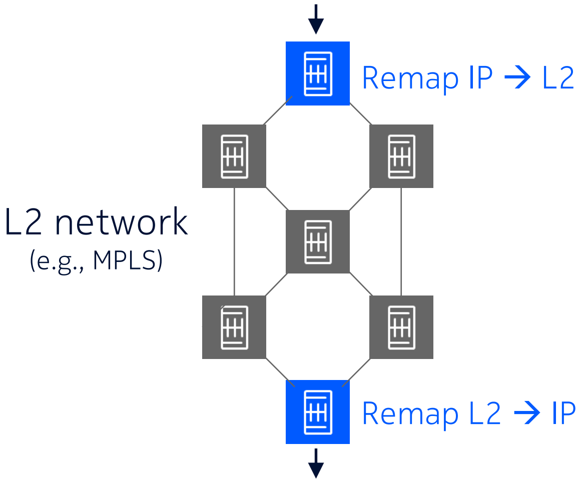

IP-L2 remapping table

Function	IP	L2 "proxy"
L4S identification (enabling prioritization)	ECN = ECT(1)	E.g.: MPLS TC* = X
L4S congestion marking	ECN = CE	E.g.: MPLS TC* = Y

* Traffic Class

→ **For prioritization, one L2-proxy value suffices**

→ **For ECN-marking, two L2-proxy values are required**

- **Same principle also applicable to other Layer-2 networks** (e.g., VLAN/.1p)
- Can also use DSCP-based proxies wherever possible
- IP-based tunneling (e.g., GTP) requires proper inner-outer IP-header ECN transfer [draft-ietf-tsvwg-ecn-encap-guidelines-22]

NOKIA BELL LABS

L4S on Nokia WiFi Beacon shows >10x peak latency decrease

- First L4S demonstration @ BBWF 2019
- Using Nokia Bell Labs' Dual-PI2 technology
- Actively used for L4S PoCs and interop testing

Latency	P50	P90	P99
No AQM	196 ms	247 ms	273 ms
L4S	**4.7 ms**	**8.6 ms**	**12.5 ms**

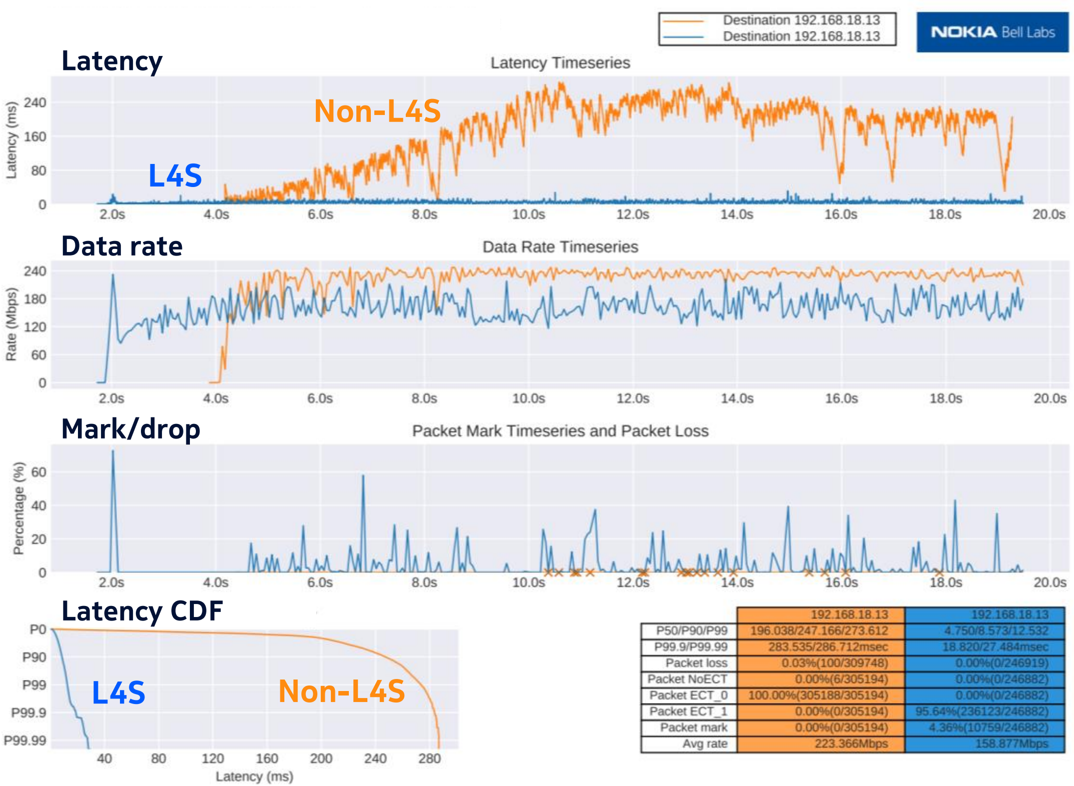

Note: Measurement on Nokia WiFi Beacon6
with reduced channel power and spectrum (20MHz)

NOKIA BELL LABS

World's first demonstration of L4S running over a fully congested end-to-end fixed network

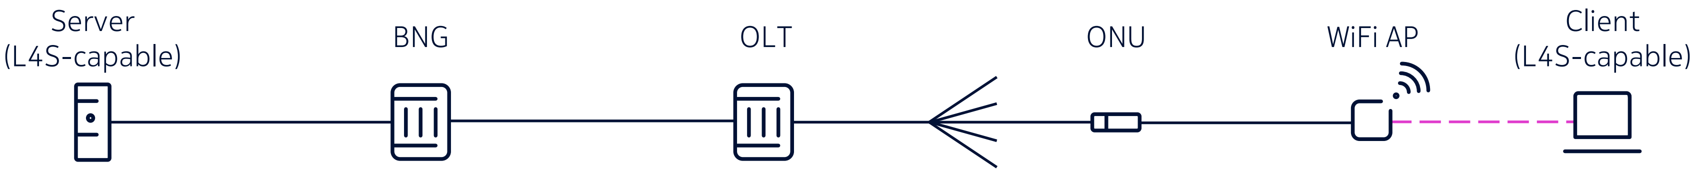

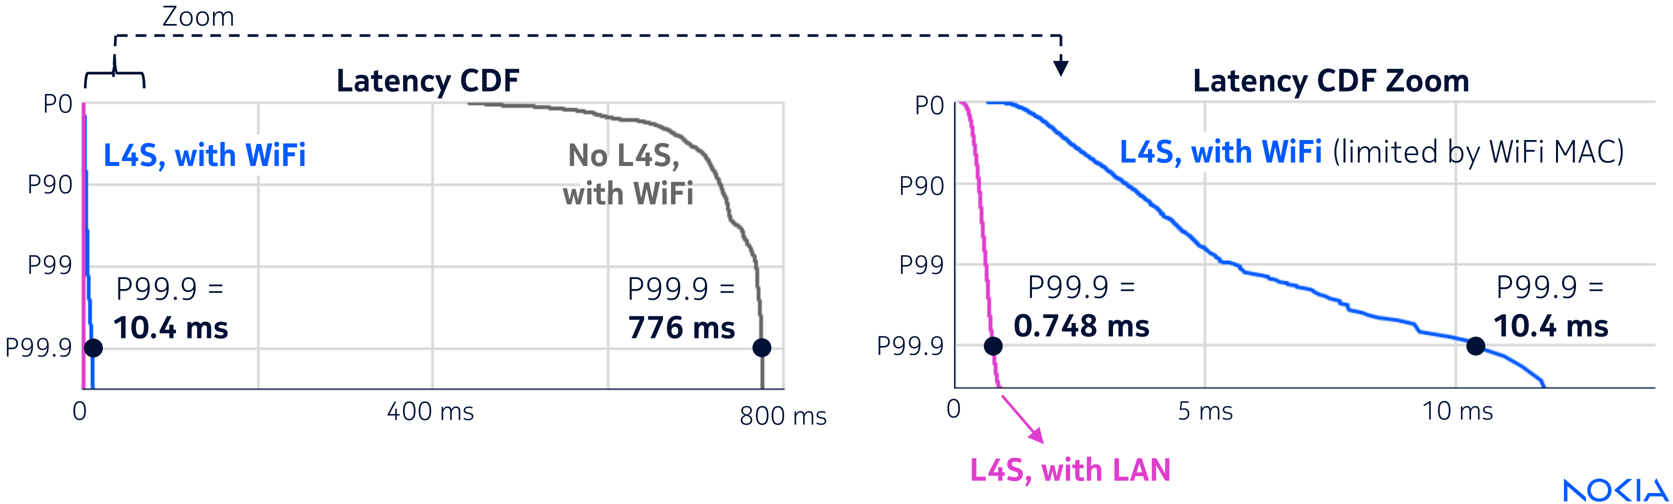

L4S allows applications to choose between 2 types of traffic

No need for the network to compromise in the middle

Classic

Buffering for **maximum throughput**

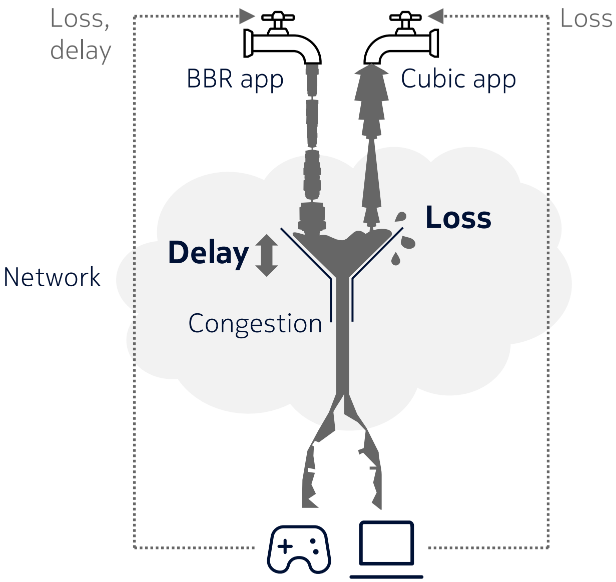

L4S

Empty buffers for **minimum latency**

[RFC9331]

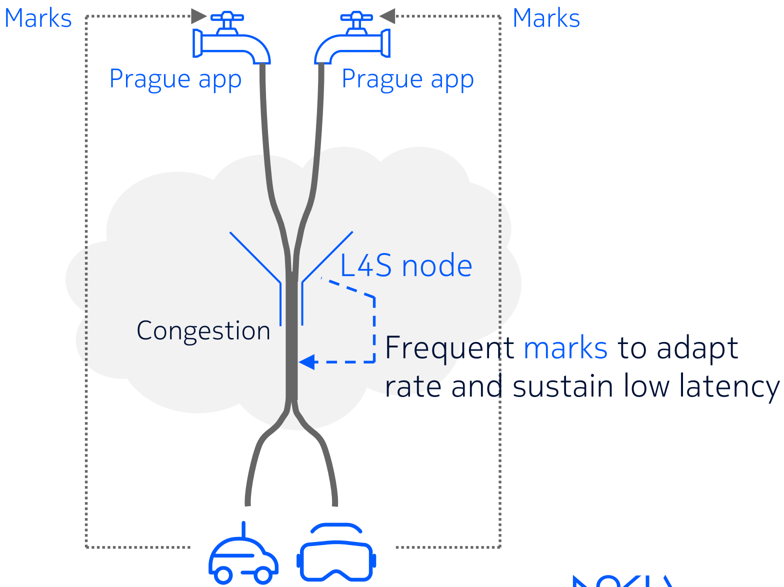

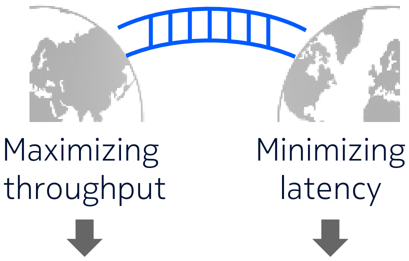

L4S allows these worlds to peacefully coexist

[RFC9332]

Further reading:

- Blog: https://bell-labs.com /l4s
- White paper: https://www.bell-labs.com/institute/white-papers/l4s-low-latency-low-loss-and-scalable-throughput/