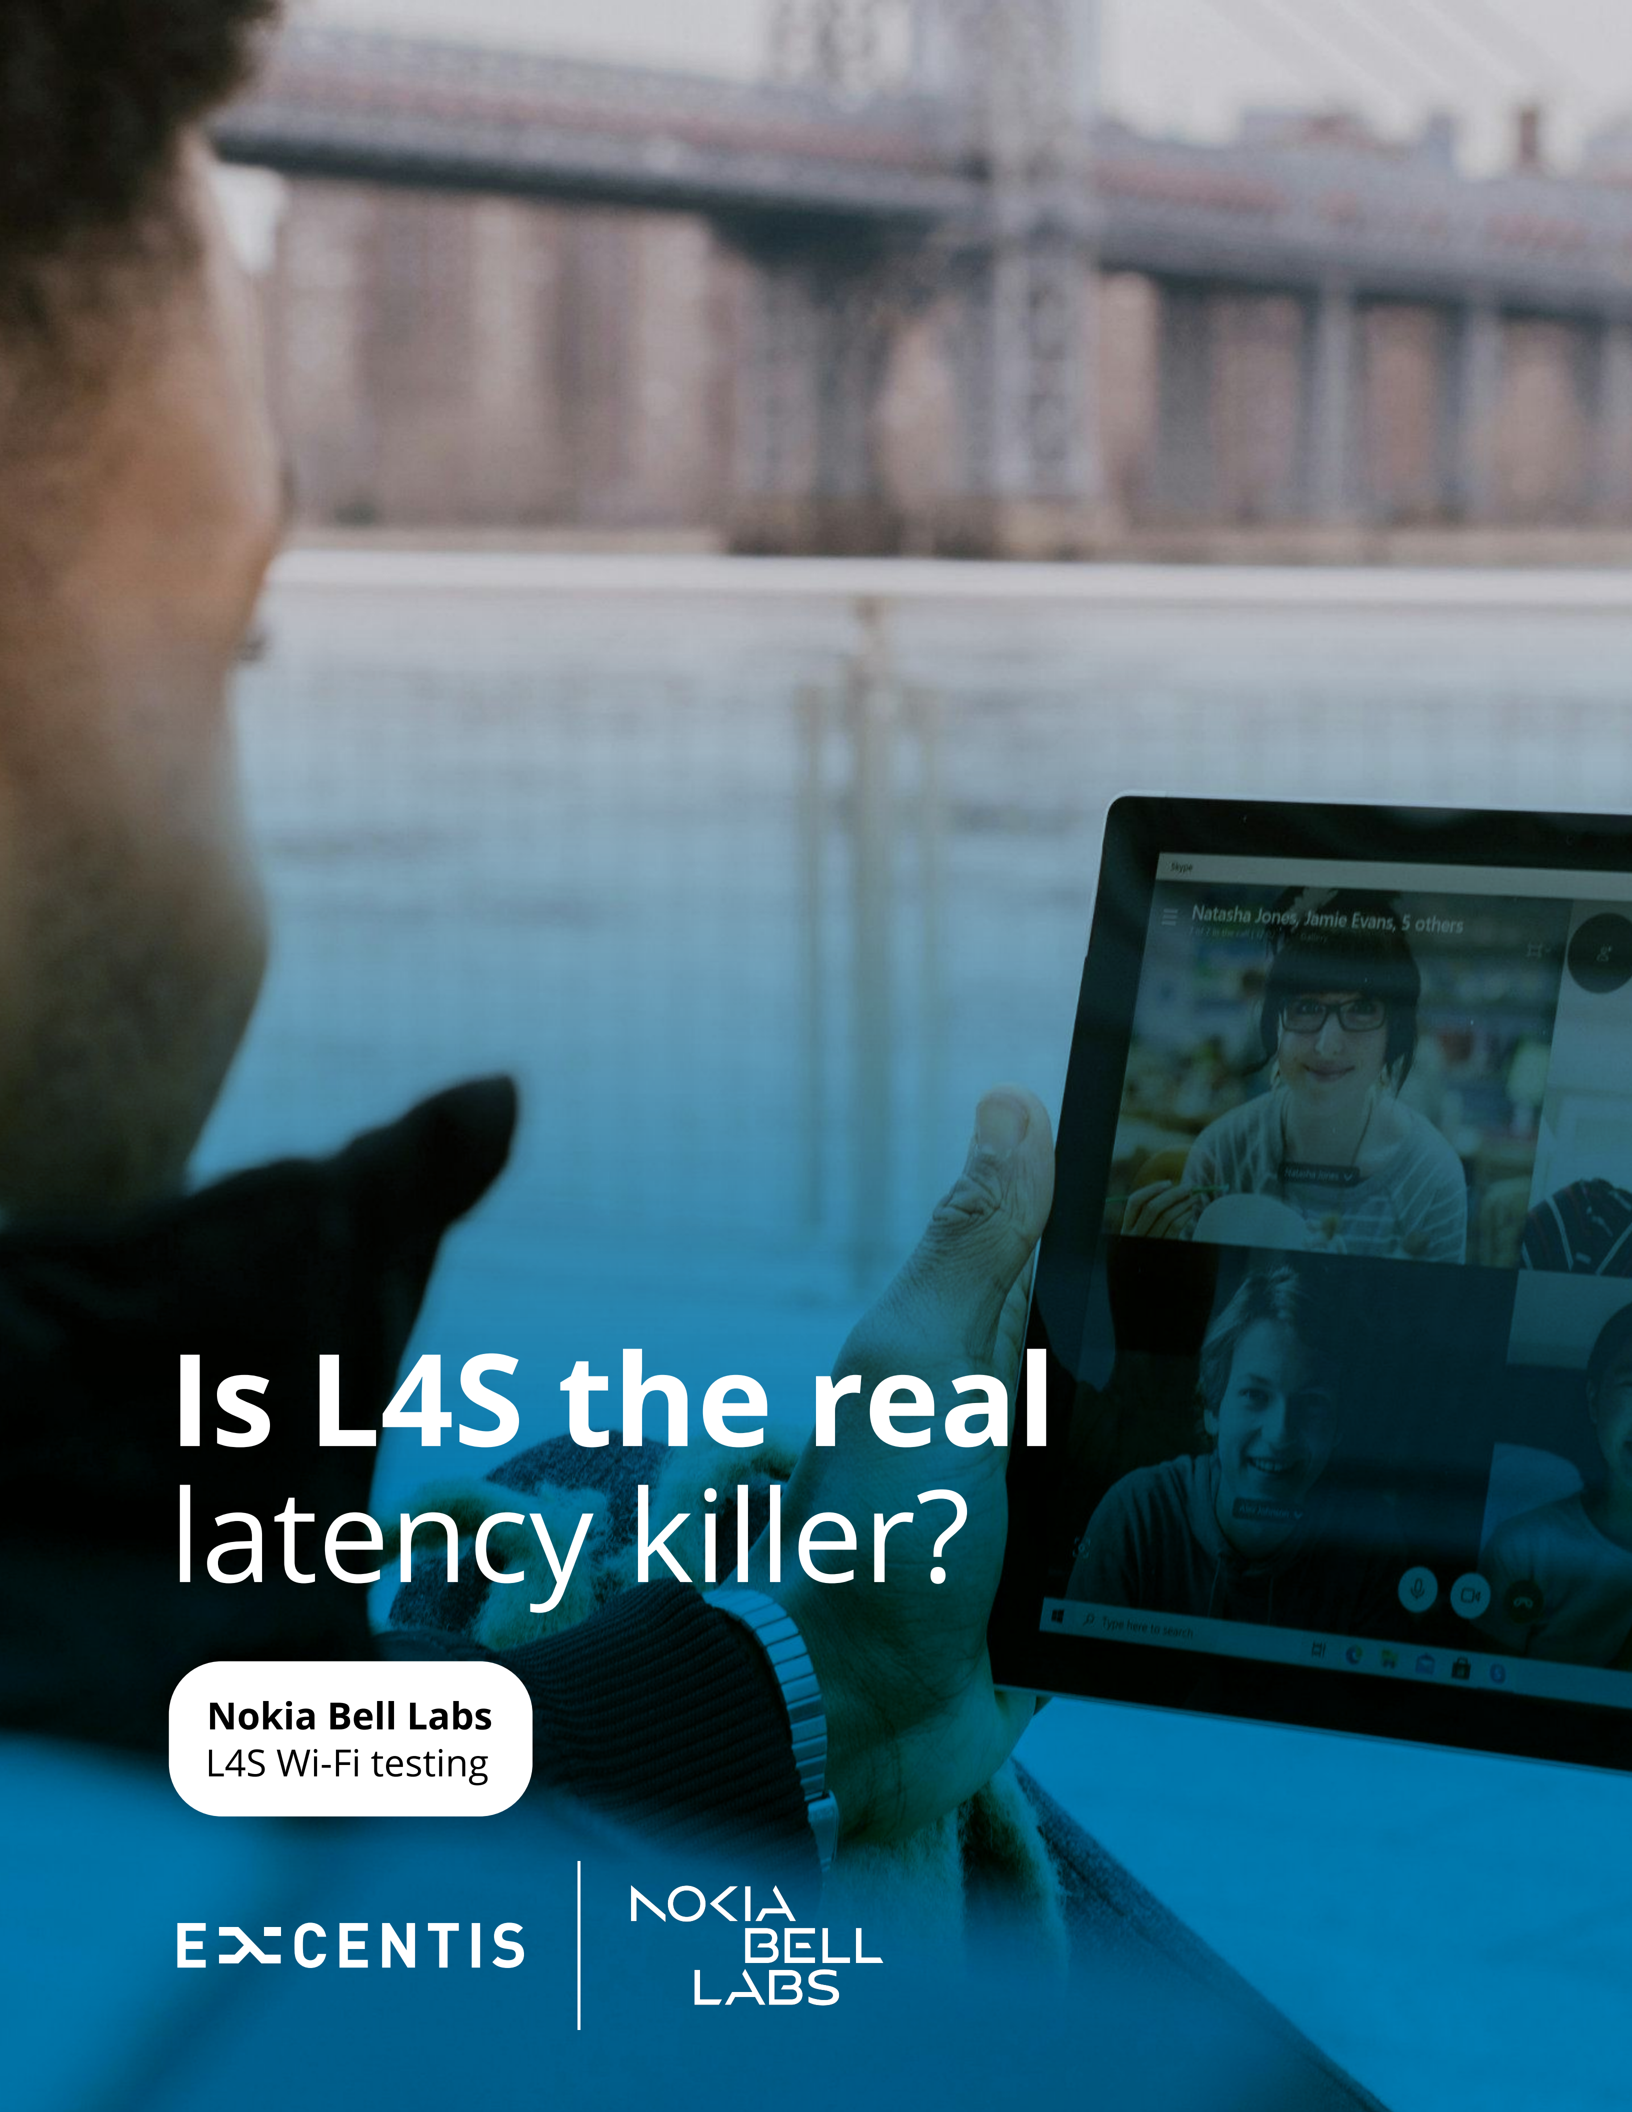

Is L4S the real latency killer?

Nokia Bell Labs
L4S Wi-Fi testing

EXCENTIS | NOKIA BELL LABS

TABLE
of Contents

INTRO

Imagine this: You're in the heat of a thrilling first-person shooter game, fully absorbed in the adrenaline-pumping action. Suddenly, your character dies without even seeing your opponent. Or picture yourself in the middle of a crucial video call, ready to secure a promotion, only to have your screen abruptly freeze. Before you get angry, consider this: The real issue might be latency.

As Internet Service Providers (ISPs) race towards innovation, they're placing more importance on low-latency services to enhance user experiences. At Excentis, we've made it our mission to tackle latency head-on. Recently, we put the latest low-latency technologies to the test, specifically L4S (Low Latency, Low Loss, and Scalable Throughput) and Wi-Fi 7, and published our findings in a white paper, "Unveil the Power of L4S and Wi-Fi 7 🧑‍💻". Expanding on our previous research, we've worked with Nokia Bell Labs to dig deeper into the impact of L4S on Wi-Fi network performance, culminating here. Together, we're diving into the world of low latency to unlock the future of network performance.

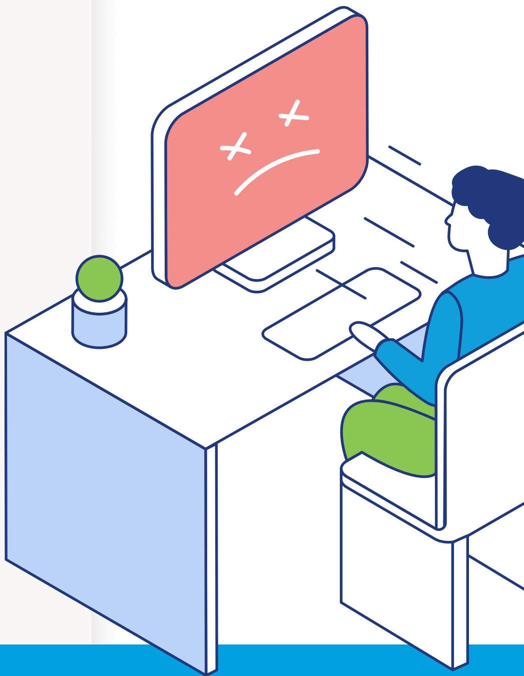

While L4S shows potential for achieving lower latency, allowing for lightning-fast transmission on networks, it also presents its fair share of challenges.

Nokia Bell Labs is leading **L4S development**, and we worked together to test their L4S-enabled access point at **Excentis Wi-Fi House.**

Are you ready to test yours?

From media acquisition delays to compatibility issues, there's a lot to consider. ISPs need to be experts in areas like Active Queue Management (AQM) and Enhanced Distributed Channel Access (EDCA) (WMM) to successfully navigate these challenges. For our independent assessment of the L4S technology in Wi-Fi, we headed to the Excentis Wi-Fi house to ensure an isolated real-world environment. Using Nokia Bell Lab's Wi-Fi 6 access point with L4S and then collaborating with a knowledgeable team of engineers from Nokia Bell Labs and Excentis, we prepared to test L4S thoroughly. We explored a range of configurations, from standard Wi-Fi 6 to L4S-optimized setups. Additionally, with a Wi-Fi 7 access point, real L4S-compatible operating systems, and the ByteBlower—our traffic generator and analyzer—we had all the tools necessary for a comprehensive evaluation.

Join us as we uncover the potential and pitfalls of L4S for Wi-Fi technology.

We'll compare it to other latency-reducing innovations like Wi-Fi 7 and fine-tune different parameters to achieve the lowest latency possible in real-world scenarios. Can these technologies play harmoniously together, or will one reign supreme? Dive into our white paper and explore the ins and outs of L4S technology and its impact on Wi-Fi networks.

CHEAT SHEET

What is L4S?

L4S, or Low Latency, Low Loss, Scalable Throughput, is a set of internet transport protocols designed to reduce latency and improve the performance of novel real-time applications in modern networks. It combines the strengths of traditional TCP and UDP to achieve lower latency at the highest throughputs, particularly benefiting latency-sensitive applications like real-time video streaming and online gaming. L4S manages signals to adapt the rate of different applications, preventing digital traffic jams and ensuring a faster, more reliable network experience.

Do all applications need to support L4S?

Cloud gaming, video conferencing, industrial control, and monitoring, and other tasks sensitive to delays benefit the most from L4S. But what about other applications? An operator, for example, has the option to classify different types of traffic into the L4S queue. For example, DNS requests, could go there too. Why? Because faster DNS means quicker loading of web pages, enhancing the overall web browsing and social media experience.

Does L4S cause the throughput to drop?

L4S ensures low latency by prioritizing its traffic, while the classic queue fills in the gaps. But what if both queues want a high throughput? No worries. The DualQ Coupled AQM ensures that both classic and L4S traffic are given equal attention, sharing the available bandwidth fairly. But is this fair scheduling enough when both queues serve similar traffic rates? Absolutely. The Dual-Queue Coupled AQM evenly splits the link's capacity, ensuring that all applications get roughly the same per-flow throughput. It calculates a probability to mark or drop packets, keeping things in balance.

Typically, with DualQ Coupled AQM, the L4S queue maintains a delay of 2ms at the 99th percentile, while the classic queue sits around 20-30ms or higher if beneficial for the performance of the link technology. And don't forget, the actual latency depends on how the AQM algorithms are configured.

Is the world ready for L4S?

Well, L4S and its AQMs handle network queueing delays, but not the underlying delays caused by things like media acquisition in Wi-Fi. To fully meet the needs of latency-sensitive applications, devices might need some tweaks. However, we're optimistic that with a simple configuration or firmware update, existing devices can become L4S-ready.

In our lab, we're putting this to the test. With Nokia Bell Lab's Wi-Fi 6 access point, we've compared various configurations, from standard Wi-Fi settings to L4S optimization. Plus, we have ByteBlower, Excentis' traffic generator and analyzer, ready to simulate L4S traffic and compare it with a Wi-Fi 7 access point that doesn't support L4S. Let's see how L4S can make a difference

Which
BLOCKS
are necessary to build the L4S technology?

When we want to make our internet faster, we usually start by fixing the bottleneck. That's like unclogging the drain to let the water flow smoothly. An L4S-ready Wi-Fi access point does just that, using fancy queueing tricks and Active Queue Management (AQM) to manage different types of traffic like a boss. **Could L4S be the internet latency superhero, swooping in to rescue the day?**

Our study? Well, it's like a grand experiment with toy blocks. We built it layer by layer, starting with basic Wi-Fi 6 setups. Then, we tapped into the power of Dual-Queue Coupled AQM (or DualQ AQM for short) to evaluate L4S performance. Finally, we dove into the MAC layer to configure latency-optimized EDCA/WMM parameters.

Nokia Bell Labs brings L4S-enabled Wi-Fi

Your Wi-Fi access point acts as the gatekeeper of your internet connection, whether at home or in a bustling public hotspot. When the crowds gather and everyone wants a piece of the Internet pie, the Wi-Fi link often becomes the bottleneck, slowing down the flow of data.

To tackle this issue, we turned to Nokia Bell Lab's cutting-edge Wi-Fi 6 access point. This allowed us to configure various setups, ranging from standard Wi-Fi 6 to the innovative L4S, following Nokia Bell Lab's recommendations for parameter optimization. L4S isn't just about speed; it's about smartly managing traffic to ensure smooth sailing for latency-sensitive applications while still being fair to all users.

How does it work? Just like in a busy intersection, instead of letting cars pile up at the stoplight, the access point uses advanced queueing mechanisms and Active Queue Management (AQM) techniques to keep the traffic flowing smoothly. One such technique, Dual-Queue Coupled AQM, ensures that urgent traffic doesn't pile up, and the optimized EDCA/WMM parameters ensure that everyone gets a very regular green light while preventing a gridlock.

With Nokia Bell Lab's Wi-Fi 6 access point featuring L4S technology and our teams working together, we're feeling the excitement—**the future of the internet is here**—promising faster speeds, lower latency, and a more equitable distribution of bandwidth for all.

Active Queue Management (AQM)

Have you ever been on a congested highway where traffic still seems to flow smoothly? DualQ AQM operates similarly on our Wi-Fi networks, thanks to intelligent technology that predicts and prevents congestion. It's as if we have two lanes on every road—one for priority traffic and the other for regular traffic—ensuring everyone reaches their destination. The priority lane experiences minimal delays, while the regular lane has just enough delay to keep the network fully optimized.

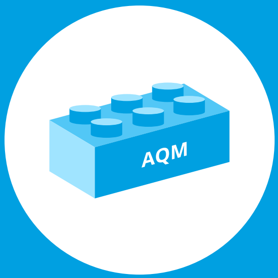

It's as if we have a road with three lanes, of which one is a HOV (High Occupancy Vehicle)-lane. Cars with more than one passenger (or priority vehicles) can use these lanes to bypass the queues formed by all single occupancy cars. At the end of the HOV-lane they merge back onto the normal road but they have bypassed the traffic jam as such experiencing much less delay.

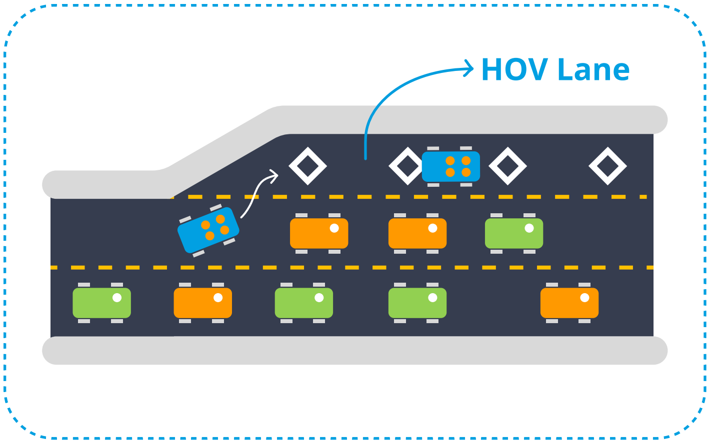

DualQ AQM uses two queues to separate and prioritize the L4S traffic from the classic traffic, ensuring very low latency for L4S and acceptable and sufficient latency for classic traffic.

If you're eager to delve into **advanced latency measurement** join our upcoming workshops on low latency.

Learn more →

Now, let's apply this to our Wi-Fi access points. They monitor traffic from multiple devices, much like cars on the road, ensuring a smooth journey for all. When traffic gets heavy, the access point buffers packets, similar to cars lining up to merge onto a busy highway. With properly managed queues, it can efficiently handle bursts of traffic, much like cars exiting a congested lane one by one.

With L4S, traffic endpoints employ intelligent algorithms to minimize delays and congestion. With DualQ AQM, priority traffic zooms into the fast lane, while regular traffic flows into the standard lane. If congestion arises, the priority lane sends signals to adjust traffic flow, mimicking how a driver receives signals from a car's GPS and slows down to avoid getting caught in a traffic jam.

Our access point handles traffic from multiple clients connecting to the internet and communicating with each other over Wi-Fi.

With data flowing in from different sources and heading to various destinations all at once, it's like managing a busy intersection where cars need to be organized and let through in an orderly fashion. Sometimes, when the outgoing network can't keep up with incoming traffic, the access point needs to slow down and queue up packets, much like cars waiting their turn at a traffic light. This queue helps manage sudden bursts of traffic, ensuring a smooth flow of data without overwhelming the system. Think of it like multiple trucks rushing towards different lanes but exiting through a single lane; they need to line up and depart one by one.

In essence, DualQ AQM ensures a seamless journey for all traffic on our Wi-Fi networks, eradicating delays and paving the way for a smoother, more efficient future.

How does the access point handle these queues?

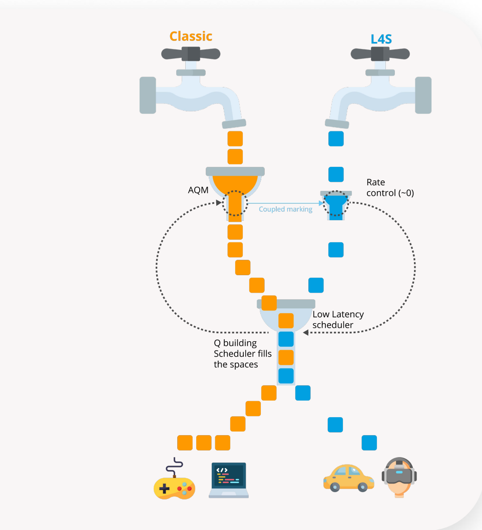

In the world of L4S, devices at both ends of the connection use smart algorithms to keep delays minimal and traffic flowing smoothly. When we combine L4S with DualQ AQM, traffic that supports this smart congestion management system is directed into the L4S queue, while traffic using traditional congestion control methods goes into the Classic queue.

The L4S queue communicates congestion using special markings known as Explicit Congestion Notification (ECN) markings in the IP packet header, allowing the receiver to signal the sender to ease up on the traffic flow. This helps prevent queues from building up on the network. Meanwhile, the classic queue uses a different method, first delaying and then gently dropping packets, to avoid congestion. Together, these queues ensure that our network runs efficiently, even during busy times, paving the way for a smoother digital experience.

Wireless Multimedia (WMM)

In a crowded room, chaos ensues if two people speak simultaneously, so everyone must wait for their turn to talk. That's like wireless devices trying to transmit data—they all compete for bandwidth, causing delays and frustration.

But fear not! By following Nokia Bell Lab's recommendations and tweaking the WMM parameters, we can bring order to the chaos. We set a high minimum and a low maximum wait time for each device, keeping it spacious and short. This means fewer collisions, less waiting around, and ultimately, lower latency. It's like giving everyone a polite short turn to speak, one after the other, making the conversation efficient and smooth, without monologues.

Enhanced Distributed Channel Access (EDCA)

The WMM parameters allow us to set the maximum wait time to a chosen limit, and in this case, we set it lower for all devices. We compensate the lower maximum with a much higher minimum, avoiding too many collisions. On top of that, we also limit how long one device can speak at a time. This greatly limits the time that any device must wait before it can retry, which significantly reduces latency.

Why does this work? In a multi-client environment, it is not a surprise to see some transmissions that don't get that acknowledgement. In our context, think of sent data that was dropped due to full buffers and hence the transmitter never receives the acknowledgement. In a congested scenario, this can happen multiple times in a row. The problem is that the transmitter doubles its wait time for each retry and this quickly results in high latencies. The maximum wait time is one of the EDCA parameters. By setting it to a lower value, we avoid the situation below.

Note that this is an experiment and not a solution that will work in every situation! A solution that works in a single-user situation will not necessarily work with five or twenty active clients.

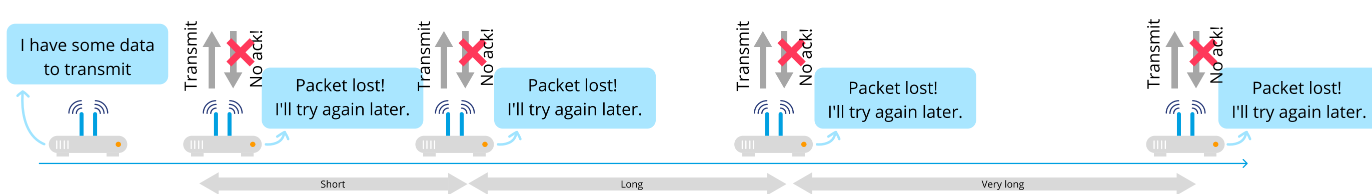

Evaluation environment &

GAME PLAN

In our quest to assess the impact of L4S in real-world scenarios, we turned to Nokia Bell Lab's Wi-Fi 6 access point and their renowned knowledge of various configurations, including L4S optimization. Additionally, we incorporated a Wi-Fi 7 access point and ByteBlower for traffic simulation.

Our testing environment is not your typical lab setting. Instead, we conducted all tests at the Excentis Wi-Fi house—a real-life dwelling situated in a remote location, devoid of external interference. This house is constructed with authentic building materials, featuring fully furnished rooms and household items that replicate a typical home environment.

By utilizing the Excentis Wi-Fi house, we ensured that our evaluations accurately replicate the conditions experienced by end-users. This approach allows us to precisely measure the impact of parameter changes and gauge the effectiveness of L4S optimization. In this immersive environment, we uncover the true potential of L4S in enhancing network performance and delivering unparalleled user experiences.

Traffic analysis

In assessing the effectiveness of L4S, we focused on TCP-based network traffic, analyzing the Round-Trip Time (RTT) between the device and the Internet. By monitoring RTT during long-duration connections and excluding initial spikes, we gained valuable insights into network performance.

By harnessing the power of L4S technology, we can revolutionize the gaming experience for users, eliminating lags and ensuring seamless gameplay in a futuristic, interconnected world.

L4S traffic is traffic where L4S ECN markings are done and where the receiver feeds back to the sender, and the sender takes actions.

Test scenarios

Our testing journey begins with a standard Wi-Fi 6 access point configuration, progressively advancing to an experimentally optimized Wi-Fi setup. At each stage, we scrutinize the improvements compared to the previous configuration. **We tested the following scenarios:**

1. **Impact of L4S**
   Wi-Fi 6 client on a Wi-Fi 6 access point with L4S enabled, compared to the standard configuration.
2. **Impact of EDCA (WMM) parameters**
   Wi-Fi 6 client on a Wi-Fi 6 access point with L4S enabled and optimized Wi-Fi configuration (low latency optimized values for EDCA/WMM).
3. **Impact of L4S vs Wi-Fi 7**
   Comparison with Wi-Fi 7 to discern differences in performance and optimization.

Case Study

Meet Alex, a real life example who loves playing a first-person shooter cloud game. He never wants to miss a shot and certainly needs to have a clear view on his enemy. He uses his gaming PC in his room to ensure that he can concentrate fully without any distractions. Because of the distance from his computer to the AP, the Wi-Fi signal strength is not optimal.

Additionally, while playing, the PC's operating system decides to download a firmware update. Both the game and the download compete for wireless resources. How can the new fascinating technologies help him?

Hey Alex!

Device placements

Lower floor

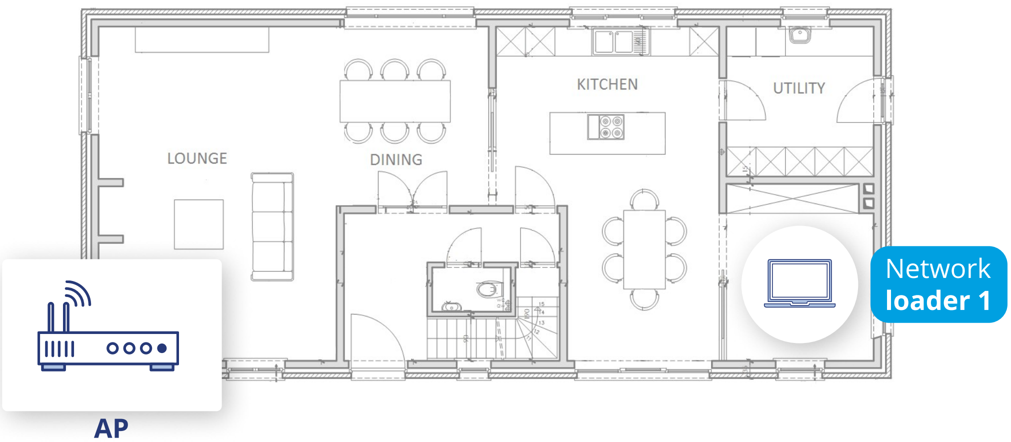

Upper floor

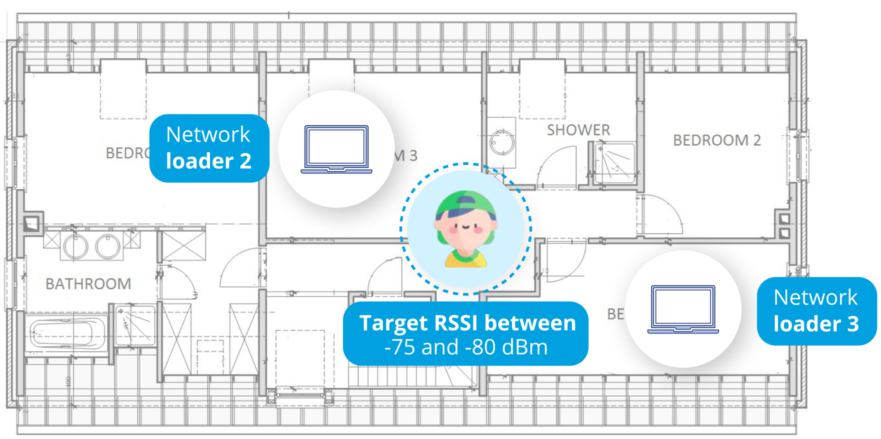

In a typical setup, the access point resides in a corner of the house, providing coverage throughout the building. However, this configuration may not offer the best connectivity for users like Alex.

L4S technology is a game-changer for users and operators alike. With L4S, there's no need for additional access points or costly installations. Excentis specializes in pinpointing latency issues and harnessing the power of L4S to streamline operations and reduce expenses.

To evaluate of the real-life behavior of the access point, we made our test setup mimic the real “typical” internet delay. In the following test cases it is called the “base delay” (+8ms).

Test
SCENARIO 1
Impact of L4S

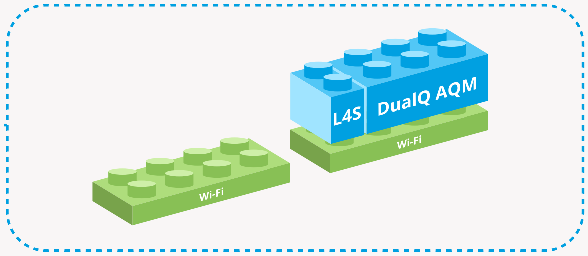

In cloud gaming, where swift reactions are crucial, maintaining high-quality video and minimal lag is paramount. With our colleagues from Nokia Bell Labs, we determined that cloud gaming applications, such as NVIDIA GeForce Now, can demand up to 40Mbps for seamless gameplay.

Cloud gaming is also an application that can scale down its throughput demand. For instance, by sacrificing video quality under highly active video scenes sent from the cloud server to the end user to avoid lag. Under these conditions, lag is very important, while video quality is not because our eyes can also only process rough but timely cues to assess the remote situation.

This all makes cloud gaming the perfect candidate for evaluating scalable congestion control in L4S on our Wi-Fi access point! We start with validating cloud gaming requirements with a single client, which happens to be congested by a firmware download running on the same laptop.

Wi-Fi 6 access point without L4S and Wi-Fi 6 client

First, we establish a benchmark for comparisons. We have a single Wi-Fi client at a distant location with low Wi-Fi signal strength. Our gamer is playing his favorite cloud game. How does it perform on an access point, without any AQM or L4S support configured at all?

Traffic Type	Average RTT Cloud gaming [ms]	Peak RTT Cloud gaming [ms]	Average Throughput Cloud gaming [Mbps]
Non-L4S	102+8	419+8	29
L4S	102+8	212+8	35

Note: The peak RTT is lower with L4S traffic, but this is not a result of L4S it just occurred that a lower peak was seen than during the non L4S test.

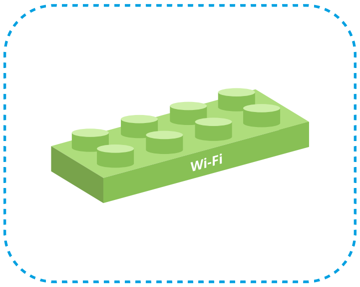

Even when the Wi-Fi client tries to use L4S, it is not supported on the access point and there are no actions taken to minimize the latency. Because the latency is so high, our users won't be happy to play their games on this link.

However, what the above results also show is that the introduction of L4S traffic on a network that does not support it has no negative impact either. This will allow a more gradual roll-out of L4S without fear of adverse effects.

Single client: L4S enabled on the Wi-Fi access point

We've now reconfigured our Wi-Fi access point to use L4S Dual-Queue Coupled AQM. We expect to see the advantages of L4S now. In a later stage, we will also tweak other parameters (EDCA) but for now, everything is left to its defaults. Let's see if the performance improves.

	Traffic Type	Average RTT Cloud gaming [ms]	Peak RTT Cloud gaming [ms]	Average Throughput Cloud gaming [Mbps]
L4S disabled on access point	Non-L4S	102+8	419+8	29
	L4S	102+8	212+8	35
L4S enabled on access point	Non-L4S	14+8	59+8	38
	L4S	8+8	45+8	35

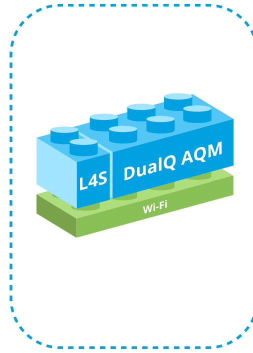

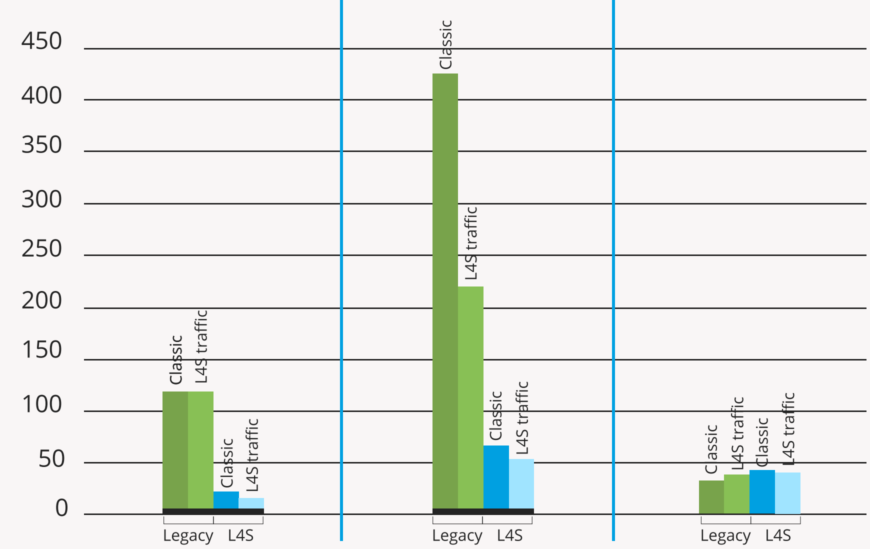

Average RTT Cloud gaming [ms] **Peak RTT Cloud gaming [ms]** **Average throughput Cloud gaming [Mbps]**

Now we start seeing some results! Excluding the 8ms base delay in the network, the L4S traffic shows us **92%** lower Wi-Fi latency than the standard Wi-Fi 6 configuration! The classic traffic also benefits from AQM and gives an improvement of about **86%** relative to the standard Wi-Fi 6 configuration.

With this L4S-aware access point, L4S traffic shows a RTT that is about 45% relative to the Classic traffic. This is achieved without loss for the L4S traffic and no significant throughput degradation. In addition to the lower latency, the jitter also was significantly reduced.

Multi-client

We continue with our gamer, who is still playing, and the OS system update is still downloading. Only this time, the Wi-Fi link is additionally congested by other family members who are using bandwidth for uploading holiday videos.

Can the L4S support of our access point again improve the gamer's user experience?

L4S on Wi-Fi 6 access point

We had acceptable latency performance when we were the only users of our Wi-Fi network. What happens if the rest of my family also wants to get a piece of the pie?

	Traffic Type	Average RTT Cloud gaming [ms]	Peak RTT Cloud gaming [ms]	Average Throughput Cloud gaming [Mbps]
One client active	Non-L4S	14+8	59+8	38
	L4S	8+8	45+8	35
Multi-client	Non-L4S	46+8	376+8	4.1
	L4S	43+8	264+8	12

Multi-client

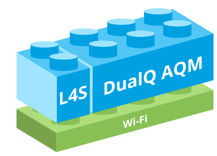

As expected, the latency increases significantly because the extra family members are using the same network resources. Multiple devices are connected now, and they all try to send data. The Wi-Fi devices simultaneously ask for transmit opportunities. This strain on the Wi-Fi connection requires retries to be allowed to transmit on the network. In addition to the latency increase, the throughput for our gamer has dropped significantly as well.

It gets even worse when we look at what happened to the download of the OS update. The OS update that was running in the background on the gamer's laptop even dropped after a few seconds!

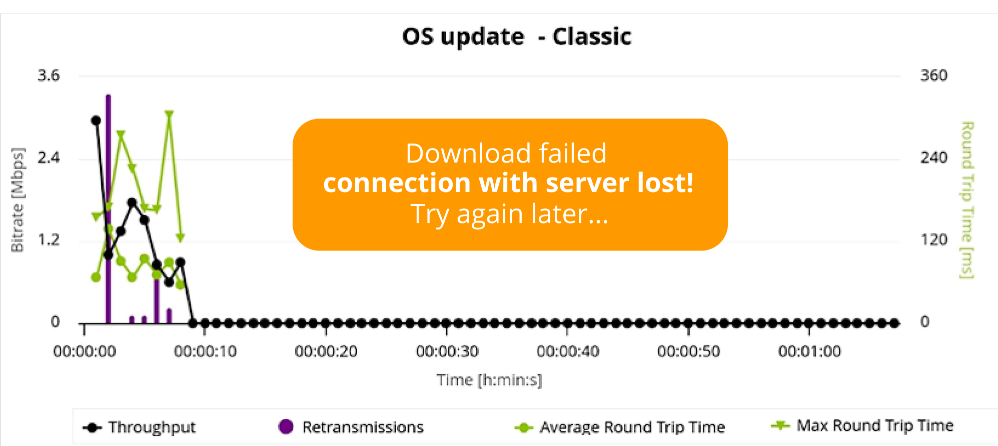

This clearly shows that even when using L4S, the delay introduced by the access network technology becomes very visible. For heavily populated Wi-Fi networks, the new bottleneck for running good performing latency-sensitive applications is the delay caused by media acquisition.

Test

SCENARIO 2

Impact of WMM parameters

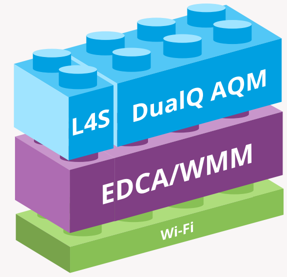

Single Client

Is this the best we can get from L4S on Wi-Fi links? Are there other parameters that we can tweak to increase performance?

The Wi-Fi specifications define the Enhanced Distributed Channel Access (EDCA) parameters for WMM, which play a role in the media acquisition of Wi-Fi devices. Limiting the maximum wait time for transmit opportunities significantly impacts latency. However, increasing the minimum wait time and reducing the sending time can lead to a reduction in the maximum throughput achievable by a single Wi-Fi client.

While this is most important for multi-client environments, the impact of changing these parameters for a single Wi-Fi client should not be ignored! To find out, we compared cloud gaming on a congested Wi-Fi link on an access point configured with:

- Wi-Fi 6 and L4S disabled
- L4S on Wi-Fi 6 using legacy WMM parameters
- L4S on Wi-Fi 6 using optimized WMM parameters

Let’s see what our tests reveal!

We ran the tests in a single Wi-Fi client on a distant location, with lower Wi-Fi strength. The test simulates playing a cloud game while its link is congested by the OS update running in the background.

L4S on Wi-Fi 6 with optimized WMM parameters

Can the access point configuration help us to further improve the latency? Nokia Bell Labs developed an improved set of Wi-Fi WMM parameters to minimize the delay introduced by media acquisition. We gave it a try!

	Traffic Type	Average RTT Cloud gaming [ms]	Peak RTT Cloud gaming [ms]	Average Throughput Cloud gaming [Mbps]
No L4S - default WMM parameters	Non-L4S	102+8	419+8	29
	L4S	102+8	212+8	35
L4S - default WMM parameters	Non-L4S	14+8	59+8	38
	L4S	8+8	45+8	35
L4S - optimized WMM parameters	Non-L4S	26+8	76+8	32
	L4S	4+8	19+8	32

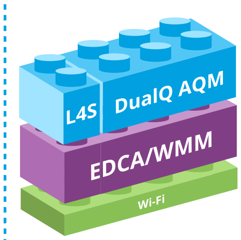

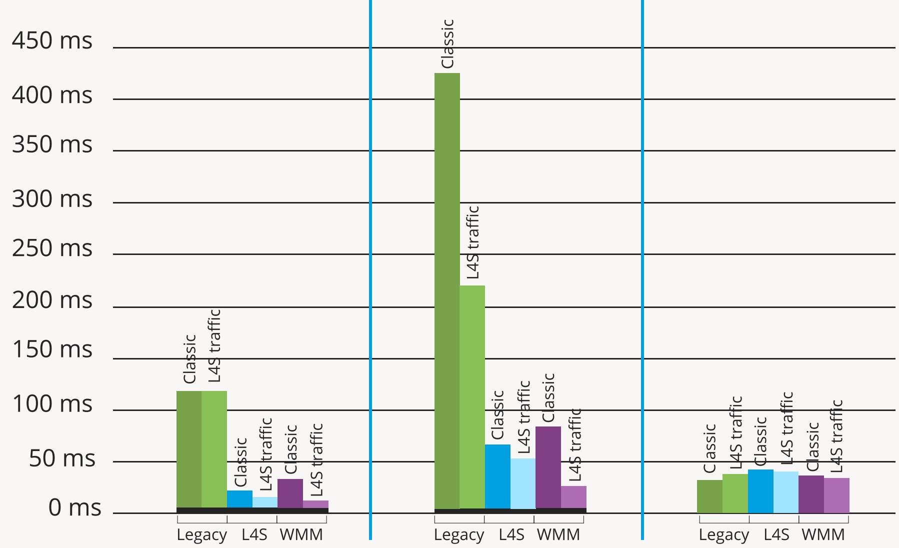

Latency on L4S traffic is much lower now (down from 8 to 4ms) while our throughput is mostly maintained! The classic traffic did not seem to benefit from the optimization attempt.

There is a trade-off though between latency and throughput when using the WMM EDCA parameters, certainly for single Wi-Fi client scenarios. Even though the OS update download is not a time-sensitive process, its throughput dropped a bit and will take a bit longer.

Multi-client

We revisit our gamer, who is still playing his game while his PC is still working on downloading OS updates. In the meantime, the Wi-Fi link is again congested by other family members who are uploading holiday videos.

With the default WMM parameters, the gamer's throughput dropped, and latency was very high. Let's see if the optimized WMM parameters helped.

L4S on Wi-Fi 6 with optimized WMM parameters

As shown before, using the legacy Wi-Fi WMM parameters did not seem to be a success. Like what we did for the single client test case, let's see what kind of improvements the access point configuration can give us.

	Traffic Type	Average RTT Cloud gaming [ms]	Peak RTT Cloud gaming [ms]	Average Throughput Cloud gaming [Mbps]
Default WMM parameters	Non-L4S	46+8	376+8	4.1
	L4S	43+8	264+8	12
Optimized WMM parameters	Non-L4S	37+8	103+8	6.8
	L4S	15+8	57+8	4.7

Multi-client

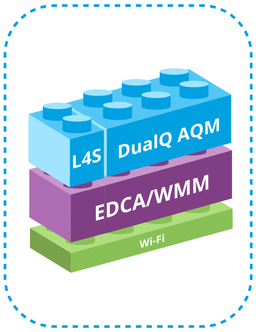

The throughput for our game is not optimal yet and the video quality will still be lowered. However, the latency has improved enough that it can make our user a happy gamer again.

And moreover, our OS update in the background could continue to work!

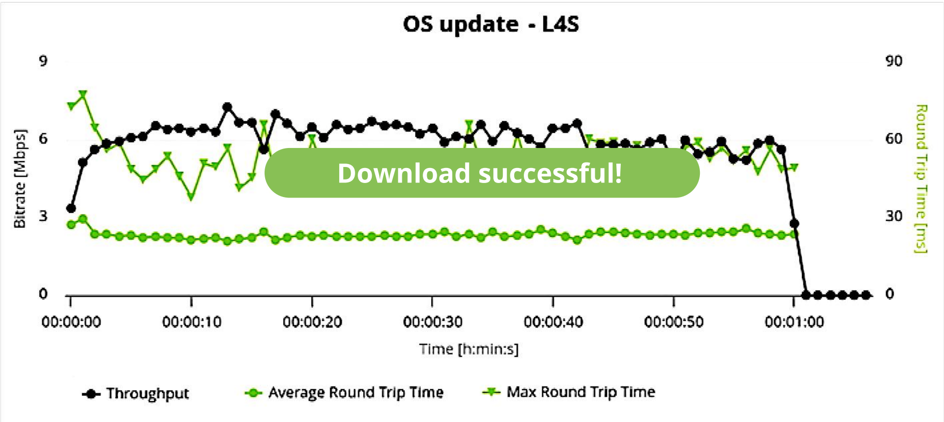

What was the influence for the gamer's family? The average total upload speed decreased from 150Mbit/s to 90Mbit/s. Again, **we see the trade-off between latency and throughput!**

Test

SCENARIO 3

Impact of L4S vs. Wi-Fi 7

Single client

Wi-Fi 7 slowly starts introducing itself on the market. L4S is not defined within the Wi-Fi 7 specification. However, Wi-Fi 7 also promises its own improvements with regards to latency. To compare the two, we recreated the same scenarios we tested with Wi-Fi 6, but this time using a Wi-Fi 7 wireless client. After that, we also installed a Wi-Fi 7 access point and tested it as well.

There is, however one important restriction: we don't yet have Wi-Fi 7 devices that support L4S! As a result, we can only compare the scenarios with classic traffic on the Wi-Fi 7 devices and not with L4S traffic.

Wi-Fi 7 on a Wi-Fi 6 access point

How does a Wi-Fi 7 client perform on a Wi-Fi 6 access point? How does this compare to a Wi-Fi 6 client on the same access point?

	Traffic Type	Average RTT Cloud gaming [ms]	Peak RTT Cloud gaming [ms]	Average Throughput Cloud gaming [Mbps]
Wi-Fi 6 client on Wi-FI 6 access point – no L4S	**Non-L4S**	102+8	419+8	29
Wi-Fi 7 client on Wi-Fi 6 access point – no L4S	**Non-L4S**	158+8	481+8	37

L4S on Wi-Fi 6 AP

Does the Dual-Queue Coupled AQM also improve the performance of a Wi-Fi 7 client? Since the access point does not know the difference with the Wi-Fi 6 client, we expect similar results, as we saw during our Cloud Gaming – Impact of L4S, Single Client, L4S is enabled on the Wi-Fi 6 access point test.

	Traffic Type	Average RTT Cloud gaming [ms]	Peak RTT Cloud gaming [ms]	Average Throughput Cloud gaming [Mbps]
Wi-Fi 7 client on Wi-Fi 6 access point – no L4S	**Non-L4S**	158+8	481+8	37
Wi-Fi 7 client on Wi-Fi 6 access point – with L4S AQM	**Non-L4S**	32+8	301+8	30
Wi-Fi 6 client on Wi-Fi 6 access point – L4S enabled	**Non-L4S**	14+8	59+8	38
	L4S	8+8	45+8	35

When we exclude the base delay (8ms), we see that the average RTT **improved by about 80%!** There are still high peaks, but the overall jitter is much lower. There was no significant decrease in throughput, but it was less stable.

Our user will be able to play his favorite cloud game again. It will be responsive, but he might still experience some hiccups and a temporary decrease in video quality.

L4S on Wi-Fi 6 with optimized WMM parameters

We also tried the Wi-Fi 7 client on the access point with optimized WMM parameters. Does it perform similarly to the Wi-Fi 6 client again? And better than with legacy WMM parameters?

		Average RTT Cloud gaming [ms]	Peak RTT Cloud gaming [ms]	Average Throughput Cloud gaming [Mbps]
Default WMM	**Non-L4S**	32+8	301+8	30
Optimized WMM	**Non-L4S**	26+8	233+8	37

The average RTT decreased again by about 20% relative to the legacy WMM parameters and about 85% relative to the standard Wi-Fi 6 configuration! There might still be hiccups in the throughput though; they are normal due to the nature of classic TCP.

Still, the RTT is higher and the throughput lower than what we measured for the Wi-Fi 6 client, but that is likely due to the slightly worse signal strength of the Wi-Fi on the client device.

Even though the device has a suboptimal Wi-Fi signal, the latency and throughput allow our user to play the cloud game in good conditions!

Wi-Fi 7 with MLO

Of course, we can't exclude the comparison with full Wi-Fi 7! We created a test setup with a Wi-Fi 7 client and access point. The access point is configured for MLO on 5GHz and 6GHz.

	Traffic Type	Average RTT Cloud gaming [ms]	Peak RTT Cloud gaming [ms]	Average Throughput Cloud gaming [Mbps]
Wi-Fi 6 client on Wi-Fi 6 access point - L4S AQM enabled on access point	Non-L4S	14+8	59+8	38
	L4S	8+8	45+8	35
Wi-Fi 7 client on Wi-Fi 6 access point – no L4S	Non-L4S	158+8	481+8	37
Wi-Fi 7 client on Wi-Fi 6 - L4S enabled	Non-L4S	32+8	301+8	30
Wi-Fi 7 with MLO	Non-L4S	46+8	256+8	35

The average RTT is a lot lower **(about 50%)** than what we saw on the Wi-Fi 6 access point, using standard Wi-Fi configuration. Also, maximum RTT and jitter are more stable. There is an overall higher throughput, likely because of MLO.
When we compare it to the Dual-Queue Coupled AQM, we do see that the latency on Wi-Fi 6 + L4S is still better than here.

Multi-client: Wi-Fi 7 with MLO

For further comparison with Wi-Fi 6 and L4S, we wanted to check how Wi-Fi 7 performs in an environment with multiple clients.

How does a Wi-Fi 7 setup behave in a similar multi-client scenario? Is it able to deliver on the promise that it decreases latency in the Wi-Fi path?

Again, the Wi-Fi link is additionally congested by family members who are using Wi-Fi bandwidth for uploading some holiday videos. We tested Wi-Fi 7 with MLO compared to the different Wi-Fi 6 scenarios we tried.

	Traffic Type	Average RTT Cloud gaming [ms]	Peak RTT Cloud gaming [ms]	Average Throughput Cloud gaming [Mbps]
Wi-Fi 6 Default WMM parameters - L4S enabled	**Non-L4S**	46+8	376+8	4.1
	L4S	43+8	264+8	12
Wi-Fi 6 Optimized WMM parameters - L4S enabled	**Non-L4S**	37+8	103+8	6.8
	L4S	15+8	57+8	4.7
Wi-Fi 7 with MLO with default WMM - no L4S	**Non-L4S**	59+8	326+8	16

The throughput of the Wi-Fi 7 network is in the same range as the tests we performed on Wi-Fi 6 using L4S AQM (DualAQM) with either legacy or optimized WMM parameters, considering that we have MLO enabled here. The RTT is slightly higher here, but also very important to note is that the jitter is higher than what we saw on Wi-Fi 6 using Dual-Queue coupled AQM. Cloud gaming will likely come with unstable performance.

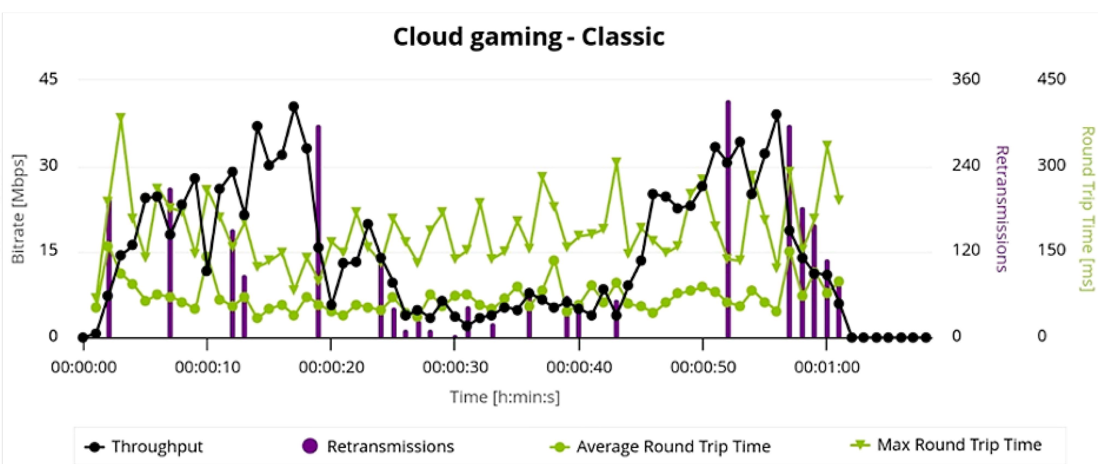

CONCLUSION

L4S works like a charm.

L4S does decrease the latency on Wi-Fi networks, but it is not a miracle solution that fixes everything without effort.

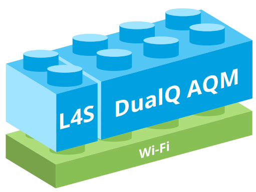

	Wi-Fi 6	Wi-Fi 6+L4S
Single client	102+8ms	8+8ms
Multi-client	152+8ms	43+8ms

With the introduction of L4S on our Wi-Fi access point, we observed significant improvements in our cloud gaming performance. **Latency improved by 70%** for multiple clients and by up to 90% for a single client. Very low latency, combined with minimal impact on throughput, creates an environment that enables responsive applications with excellent media quality.

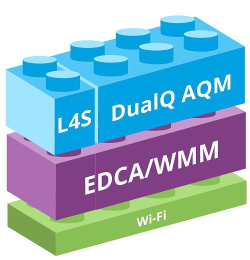

	Wi-Fi 6 + L4S + Default EDCA/ WMM	Wi-Fi 6+ L4S + Optimized EDCA/ WMM
Single client	8+8ms	4+8ms
Multi-client	43+8ms	15+8ms

With L4S enabled, it becomes clear that delays introduced by specific access technologies can no longer be ignored. The media acquisition delay becomes the dominant contributor to latency. When we want to deploy real low-latency Wi-Fi networks, **device vendors might consider new configuration parameters for their access points.**

We discovered that seemingly harmless configuration of Wi-Fi client settings can cause high latency peaks from time to time. **The software for these clients needs adjustments to minimize the impact of these features.** Having 5 seconds of high latency every 15 seconds is not what we expect when we want to achieve low latency in reduced Wi-Fi conditions!

	Wi-Fi 6 + L4S	Wi-Fi 7
Single client	8+8ms	46+8ms
Multi-client	43+8ms	59+8ms

Challenges with L4S

For L4S to truly shine, everyone needs to be on board. While it boasts impressive capabilities, our experiments revealed some hurdles that hinder its effectiveness. Here are two key issues we encountered:

Wi-Fi Parameters:

We configured specific Wi-Fi EDCA parameters for WMM to optimize low latency for all connected devices. However, some devices ignored these settings and relied on their own hardcoded configurations. This led to uneven latency distribution across stations, undermining the benefits of the optimized EDCA parameters.

Wi-Fi Scanning:

During Wi-Fi scanning—when devices search for available networks—transmission of data is halted. In some cases, these scanning intervals occurred every 30 seconds, causing significant delays or even complete traffic stoppages. This periodic disruption resulted in latency spikes that L4S couldn't mitigate.

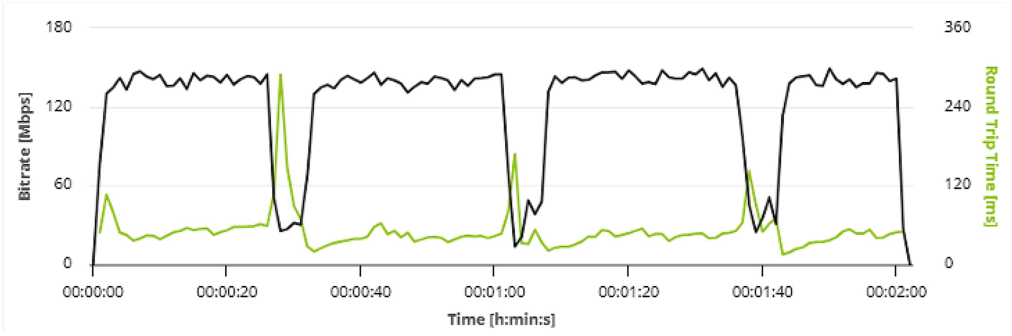

Periodical gaps in forwarding due to Wi-Fi scanning

In locations with weak signal strength, devices intensified their scanning efforts, exacerbating the latency issue even further.

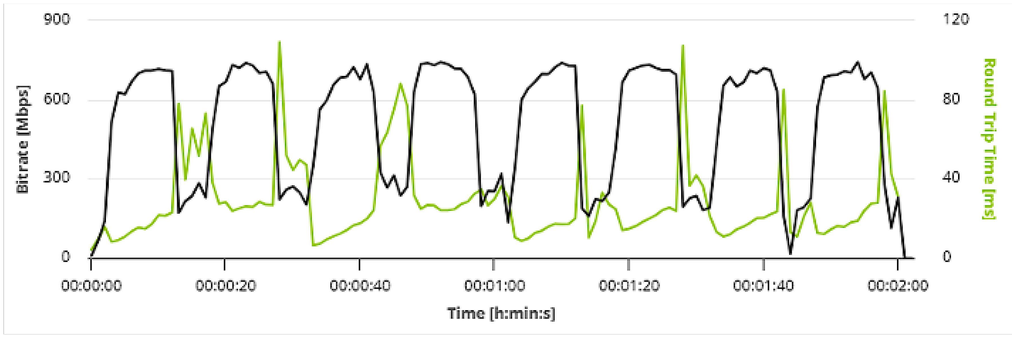

Intensive Wi-Fi scanning with tremendous latency impact

These findings highlight a critical gap in the design of devices, operating systems, and services—they often overlook latency considerations.

To deliver seamless gaming experiences and uninterrupted video calls, there must be a concerted effort to prioritize latency testing in the development process. Only then can we achieve the next-generation user experiences we envision.

Key takeaways

1. **L4S enhances latency-sensitive applications:** Implementing L4S can significantly improve the performance of latency-sensitive applications, ensuring smoother user experiences.

2. **Access Networks Influence Latency:** Each access network introduces its own latency challenges.

3. **Embrace New Technologies:** Stay ahead with innovations like Wi-Fi 7 and L4S with DualQ AQM. Discover the ideal combination for your specific use cases with Excentis' expertise.

Delve into advanced network optimization to achieve ultra-low latency performance!

Learn more here

4. **Address Flaws for Seamless Performance:** Even the most optimal technology combination may have shortcomings. Mitigate these flaws and minimize their impact on end-users through thorough testing and validation.

Test with us

What can EXCENTIS do for you?

Let's reach
new levels of low latency
together.

To achieve low latency in networking, it takes expertise and dedication. Whether you're working with Wi-Fi 6 or the latest Wi-Fi 7 technology, Excentis is here to help you reach new levels of low latency. This means better user experiences, less customer churn, and big cost savings.

What makes us different?

- Infrastructure
- Deep testing
- Independent
- Fast
- We've got the Brains
- Broad overview

Excentis and Nokia Bell Labs have laid out a detailed plan for achieving low latency. It involves understanding the network, operating systems, and devices inside out. We fine-tune everything using our knowledge of technologies like AQM, L4S, EDCA parameters for Wi-Fi WMM, and Wi-Fi 7 MLO.

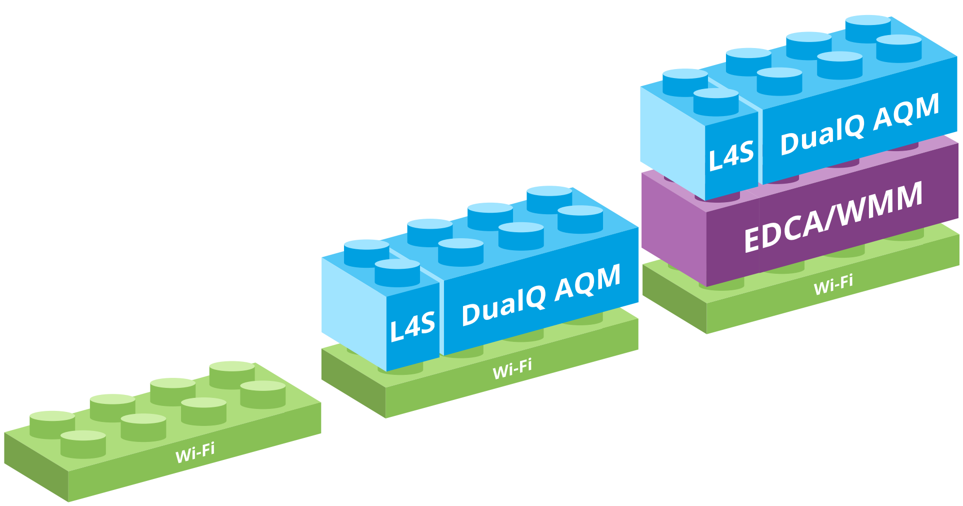

But imagine if these settings weren't tested or set up correctly. It's like building a tower with unstable blocks—your whole structure could collapse, making low latencies impossible to achieve.

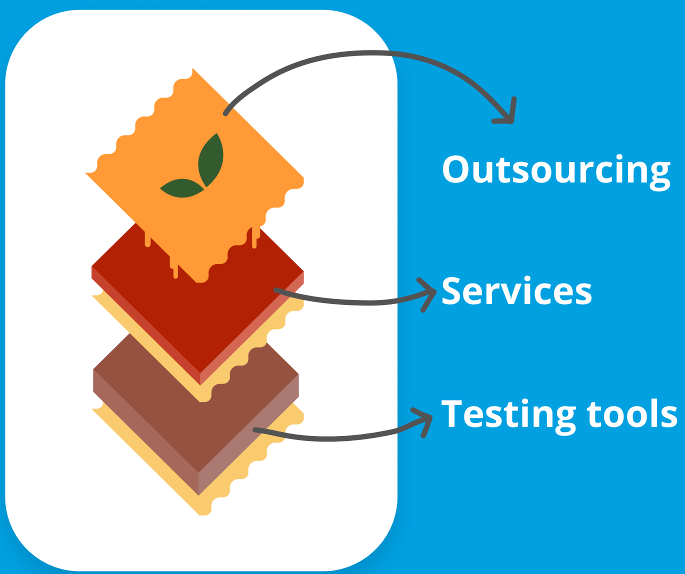

EXCENTIS lasagna?

That's where **Excentis** comes in. Our structure is like layers of lasagna. We specialize in helping you build a strong framework to optimize latency. Our capabilities range from our testing tools to complete lab outsourcing. We focus on long-term solutions to avoid problems like poor user experiences, high costs, and customer churn.

In this recent case study, we worked with Nokia Bell Labs and their L4S-enabled AP, the Nokia Beacon 6. Using our ByteBlower tool, we generated traffic and analyzed the results, all within the real-life environment of our Wi-Fi House. This setup ensures that our findings are as realistic as possible.

Our team of skilled engineers, along with Nokia Bell Lab's, conducted these tests meticulously, identifying any issues and providing actionable insights for achieving optimal performance.

With Excentis by your side, you can be confident in delivering the low-latency networking experience your users demand.

We successfully reduced initial latency

↓70% Multi-client

↓90% Single client

Optimize your latency →

We're equipped to delve deeper into your latency concerns and uncover the root causes. Through creating various scenarios and analyzing the role of each device involved, we systematically narrow down the possibilities until we pinpoint the exact cause.

Moreover, with L4S operating effectively, we're here to help you highlight its performance to all your stakeholders! Allow us to showcase the before and after latency picture in your network.

REFERENCES

RFC 9330 – Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture (ietf.org)
RFC 9331 – The Explicit Congestion Notification (ECN) Protocol for Low Latency, Low Loss, and Scalable Throughput (L4S) (ietf.org)
RFC 9332 – Dual-Queue Coupled Active Queue Management (AQM) for Low Latency, Low Loss, and Scalable Throughput (L4S) (ietf.org)
Bufferbloat – Wikipedia
RFC 7567 – IETF Recommendations Regarding Active Queue Management
IEEE 802.11e-2005 – Wikipedia
IEEE 802.11e EDCA Networks: Modeling, Differentiation and Optimization | IEEE Journals & Magazine | IEEE Xplore

2024